MULTI-SCALE PATH PLANNING FOR REDUCED ENVIRONMENTAL IMPACT OF AVIATION

BY

SCOT EDWARD CAMPBELL

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Aerospace Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2010

Urbana, Illinois

Doctoral Committee:

Professor Mike Bragg, Chair and Director of Research
Professor Victoria Coverstone
Assistant Professor Cedric Langbort
Dr. Natasha Neogi, National Institute of Aerospace
Professor Don Wuebbles

Abstract

A future air traffic management system capable of rerouting aircraft trajectories in real-time in response to transient and evolving events would result in increased aircraft efficiency, better utilization of the airspace, and decreased environmental impact. Mixed-integer linear programming (MILP) is used within a receding horizon framework to form aircraft trajectories which mitigate persistent contrail formation, avoid areas of convective weather, and seek a minimum fuel solution. Areas conducive to persistent contrail formation and areas of convective weather occur at disparate temporal and spatial scales, and thereby require the receding horizon controller to be adaptable to multi-scale events. In response, a novel adaptable receding horizon controller was developed to account for multi-scale disturbances, as well as generate trajectories using both a penalty function approach for obstacle penetration and hard obstacle avoidance constraints. A realistic aircraft fuel burn model based on aircraft data and engine performance simulations is used to form the cost function in the MILP optimization.

The performance of the receding horizon algorithm is tested through simulation. A scalability analysis of the algorithm is conducted to ensure the tractability of the path planner. The adaptable receding horizon algorithm is shown to successfully negotiate multi-scale environments with performance exceeding static receding horizon solutions. The path planner is applied to realistic scenarios involving real atmospheric data. A single flight example for persistent contrail mitigation shows that fuel burn increases 1.48% when approximately 50% of persistent contrails are avoided, but 6.19% when

100% of persistent contrails are avoided. Persistent contrail mitigating trajectories are generated for multiple days of data, and the research shows that 58% of persistent contrails are avoided with a 0.48% increase in fuel consumption when averaged over a year.

To Mom and Dad

Acknowledgements

The success of this research was dependent on many people other than myself. First, I would like to thank my advisors, Professors Mike Bragg and Natasha Neogi. Professor Bragg gave me invaluable feedback on the overall scope of the project, helped model the system, and gave me the freedom to explore a research topic outside of his central research interests. Professor Neogi provided outstanding technical guidance on the topics of controls, optimization, and air traffic systems. My graduate student colleagues Miles, Heather, Andres, Daniel, and Andy were a network of support not only in classes, but also in preparation for numerous conferences and oral exams. Thanks to my office and lab-mates Ed, Jason, Andy, Greg, and Yong – you engendered an atmosphere that made research fun, and for that I am very grateful. A special thanks to Staci, Lori, Angie, Kendra, and Barb who always answered my administrative questions and made sure that I was on track to graduate. Last but not least I would like to thank my friends and family. Mom, Dad, Jenny, Michelle, Mike, and Glenn – you all gave me a sanctuary of sanity that allowed me to unwind, enjoy life, de-stress, and prepare for the finish.

This work was funded by a University of Illinois Critical Research Initiatives grant and a NASA Graduate Student Research Program fellowship.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Background

Air transportation is critical to the world's infrastructure, economy, and quality of life. The impact of aviation on the economy is practically immeasurable due to the many indirect and direct influences that affect day-to-day life. Now more than ever, households and businesses rely on the advantages and cost effectiveness of air transport. Commercial aviation is used to facilitate global commerce, transport goods and people, and is responsible for countless jobs. In 2007, civil aviation accounted for approximately 12 million jobs, $1.3 trillion of total economic activity, and made up roughly 5.6 percent of the U.S. Gross Domestic Product (GDP) [1]. The continued advancement of the global economy is intimately tied to the well being of the air transportation system.

Aside from the economic benefits, aviation negatively affects quality of life through noise and air pollution, land use, and climate change. Emissions from airport arrivals and departures include particulate matter (PM), CO, NOx, and SOx, and a correlation between these pollutants and poor health around airports has been observed [2]. Additionally, aircraft noise has been linked to a number of health effects including hearing loss and sleep impairment [3]. The effect of aviation on climate is caused primarily from the emissions of CO_2, NOx, and H_2O [3], which are a function of the time of flight and the rate of fuel burn.

Increasing demand for air travel has pushed the utilization of the air traffic system to its limits, manifesting itself in a loss of efficiency, degradation of flight safety, and an increase in environmental impact. In addition, increasing fuel prices are straining the current economic stability of the air transportation industry. With a projected threefold increase in air traffic by 2025 [4], it is crucial to improve the current air traffic management capabilities. If the air traffic system is not modernized, there is a risk of both an economic slowdown and an increase in environmental harm due to more delays, fuel consumption, and congestion. This dissertation encompasses the development and application of a path planning algorithm to guide aircraft from departure to destination with the objective to minimize environmental impact and maximize fuel efficiency.

1.1.1 Aviation and the Environment

Aviation affects the environment through many different pathways including land use, noise pollution, local air quality, and climate. Currently, its effect on climate is not fully understood, but there are concerns that aircraft emissions might play a larger role in future global climate change than originally expected [5, 6]. The three largest aviation emissions effectors on the climate are as follows: direct emission of greenhouse gases such as CO_2, emissions of NOx, and persistent contrails [5, 6]. In general, persistent contrails are formed when an aircraft passes through an ice-supersaturated region in the atmosphere. Although the absolute consequence of persistent contrails on the environment is not known, there is evidence to suggest an impact exists [7], and it is predicted that persistent contrails have a three to four times greater effect on the climate than CO_2 emissions [8].

According to the EPA Aircraft Contrails Factsheet, contrails are line-shaped clouds composed of ice crystals that typically form in the upper atmosphere behind jet aircraft engines [9]. The lifespan of contrails is short when they are formed in dry air; however, when formed in an area of the atmosphere where the air is supersaturated with respect to ice, the contrail will persist and grow in terms of its lateral coverage [10]. The growth of contrails form a layer of upper atmosphere clouds identified as contrail cirrus. Line-shaped contrails and contrail cirrus are composed of ice crystals that reflect shortwave solar radiation and trap longwave radiation [11]. The net effect of contrails is thought to be a climate-warming effect due to the fact that longwave warming dominates over shortwave cooling as a result of the small optical thickness of most contrails observed in field measurements [12]. The global impact of contrails is estimated to be 2-8% of the global radiative forcing [12]. However, it should be understood that the effect is greater in areas of high density air traffic.

Another impact on the environment from aviation is increased emissions due to the routing of aircraft around thunderstorms. Thunderstorms are formed when an air mass becomes so unstable that convection occurs. Convective weather is a disturbance to the NAS that has significant consequences not only in terms of safety and on-time performance [13], but also environmental impact. Additionally, the accuracy and reliability of long-term forecast models for thunderstorms are not good, which cause air traffic managers to make conservative decisions in their routing of air traffic. These conservative decisions result in flight delays, including rerouting aircraft, ground holding, and en-route holding; all of which correspond to increased fuel burn, and in turn, more emissions [14]. The research described in this dissertation presents a path planning algorithm that efficiently routes aircraft around thunderstorms and areas of ice supersaturation to minimize both fuel consumption and persistent contrail formation.

1.1.2 Path Planning for Air Traffic Management

With the current capabilities of the Air Traffic Control system reaching its limits, recent research has focused on developing strategies to relieve some of the pressures air traffic controllers are facing. Further research has focused on providing pilots and air traffic controllers more information on optimal flight routes around disturbances which have the potential to decrease safety and increase the environmental impact of aviation. Continuous Descent Approaches (CDA) have been proposed to reduce noise and emissions around airports, as well as decrease fuel consumption [15]. En-route traffic optimization has been investigated from a collision avoidance perspective [16], where the cost is a function of both fuel consumption and system throughput. The remainder of this section will discuss the specifics of path planning research in persistent contrail mitigation and convective weather avoidance.

1.1.2.1 Persistent Contrail Mitigation

Persistent contrails are formed when an aircraft creates a contrail in an ice super saturated region. Airliners fly an average of 15% of their time in areas of ice supersaturation [17]. Therefore, it is prudent to devise methods to avoid persistent contrail formation, thereby mitigating the corresponding environmental effects. The avoidance of persistent contrails is a relatively broad research topic, with strategies including novel engine designs [18], fuel additives [19], and aircraft design [20].

Operational strategies for persistent contrail avoidance have also been proposed. Klima showed that heuristic contrail mitigation strategies can reduce persistent contrail coverage [21]. For the case of individual rerouting (each aircraft is rerouted independently), persistent contrails were reduced 65%-80%. For the case of weekly rerouting (custom routes are changed on a weekly basis), persistent contrails were reduced 40%-75%. The contrail mitigation reroutes computed by Klima were computed at the flight condition for minimum drag. Therefore, to accurately assess the cost/benefit trade-off the reroutes were compared to the route that minimized fuel burn (regardless of contrail formation). This comparison resulted in a 55%-85% reduction in persistent contrails and a 1%-2% increase in operational costs. Other strategies proposed by Klima include routing aircraft away from the humid tropopause, flying routes with less fuel burn, and choosing a more northerly route for transatlantic flights [21].

Because regions of ice supersaturation are very thin, it has been proposed that altitude change is the best method to avoid contrail production. Mannstein et al. showed that small changes in aircraft altitude can significantly reduce the impact of contrails [22]. Williams and Noland also assessed the viability of altitude changes on contrail formation [23]. They found that restricting cruise altitudes reduced the amount of persistent contrails formed, but with a severe fuel consumption cost. Fichter et al. found that contrail coverage could be reduced approximately 45% by flying 6000 feet lower on average with a 6% penalty in fuel burn [24]. Campbell et al. used path optimization techniques and found that persistent contrails can be avoided during a specific flight for a 2.76% increase in fuel consumption [25], or 1.48% if only 50% of the persistent contrails are avoided [26]. This dissertation will assess the viability of these methods using optimization techniques based on a realistic aircraft fuel burn model and atmospheric data.

1.1.2.2 Convective Weather Avoidance

Convective weather avoidance research has primarily concentrated on improving decision-support tools to aid air traffic controllers in making weather related reroutes. Examples of such research include work done by Kuhn and DeLaura [27,28]. An area with increasing research interest is in the development of algorithms to aid in autonomous avoidance of thunderstorms. Krozel developed an algorithm to minimize route exposure to hazardous weather by taking into account constraints on turn rate and aircraft dynamics [29]. Pannequin et. al used a model predictive control approach to find optimal trajectories around areas of static weather [30]. Their approach was formulated as an optimal control problem and included collision avoidance in addition to weather avoidance. Recently, Nguyen et. al used a heuristic approach to search the state space for solution paths around convective weather that minimize heading change, altitude change and distance traveled [31].

1.1.3 Path Planning and Optimization Techniques

Path planning has been studied in many disciplines including robotics, watercraft, aircraft, and spacecraft [32-34]. Problems involving autonomous navigation around a field of obstacles have been formulated as optimal control problems [35], potential field problems [36], and numerical approaches such as pseudospectral methods [37] and nonlinear trajectory generator [38] have been studied. Receding horizon control is another technique that has been employed in autonomous path planning research [39]. Traditionally used in the field of process control, receding horizon controllers solve a trajectory optimization problem using different levels of resolution for the near and far terms. In general, this reduces the computational burden of some optimization problems by compromising overall optimality.

Mixed-integer linear programming (MILP) has been used as a technique in the robotics and autonomous navigation fields to find the optimal path through a field of obstacles [40-45]. In this approach, obstacles are included as a set of constraints, where each constraint defines an edge of an obstacle. Each obstacle requires a set of binary variables for each obstacle edge for each time step, so the size of the problem must be controlled for reasonable computation performance. Unfortunately, MILP does not scale well with the number of binary variables in the problem formulation, and it has been proved to be NP-hard in the number of binary variables [46, 47]. Several methods have been developed to mitigate the computational struggles of MILP optimization. Earl and D’Andrea proposed an iterative algorithm to adjust the optimal route as it encountered obstacles [48]. Vitus et al. introduced a method called Tunnel- MILP which breaks the problem into three steps that are simpler to solve than the complete problem

[49]. Bellingham and Richards formulated the MILP in a receding horizon fashion, which effectively reduced the size of the problem by neglecting a detailed optimization around obstacles far from the aircraft [50]. Receding horizon MILP has been shown experimentally to guide an aircraft around obstacles in an online fashion [51]. Receding horizon MILP forms the foundation of the path planning algorithm developed in this dissertation.

1.2 Dissertation Outline and Summary of Contributions

This dissertation presents a path planning framework to deal with trajectory generation problems in the presence of multi-scale obstacles. The problem of interest is the environmental impact of aviation, and the proposed algorithm generates an aircraft trajectory which minimizes fuel burn while mitigating persistent contrail production, seeks the most optimal route around thunderstorms, and avoids collisions with other aircraft. The dissertation is organized as follows, with the major contributions of each chapter highlighted:

- **Chapter 2** introduces the atmospheric data used to model the environment in the path planning chapters of the dissertation. It also discusses contrail formation, and provides details on the source of the relative humidity data as well as the persistent contrail formation model. Convective weather is addressed using an explanation of how the radar data is obtained and how it is used in the thunderstorm model.

- **Chapter 3** shows the development of a three-dimensional path planning problem. Much of this path planning algorithm follows from existing work [40, 41]; however, this chapter contributes a much more realistic aircraft performance and fuel burn model than previously employed, and analyzes the difference between quadratic and linear cost functions.

- **Chapter 4** presents a novel cost-to-go formulation that allows for the inclusion of soft avoidance constraints in the receding horizon MILP. This chapter steps through the development of the algorithm and provides simulation results for an example scenario. An analysis of the algorithm is conducted to describe how the running time increases with problem size.

- **Chapter 5** gives an analysis of how the receding horizon controller is sensitive to disturbances of multiple scales in the environment. It looks at the effect of changing planning horizon size, execution horizon size, and time step size on the overall trajectory performance in the presence of obstacles that are of different scales, and have both hard and soft avoidance constraints. Lastly, a “best” receding horizon strategy is proposed, which can successfully adapt to a problem with multiple scales in the environment.

- **Chapter 6** shows the practical applications of this algorithm. The performance of the algorithm is tested as a technique to mitigate persistent contrail formation, and the results are compared to existing contrail mitigation strategies. Additionally, it provides simulation results for multi-scale scenarios involving both contrail mitigation and thunderstorm avoidance. Lastly, this chapter presents other practical applications for this work including turbulence avoidance, sonic boom mitigation, and aircraft icing avoidance.

- **Chapter 7** gives a summary of this work, provides conclusions, and presents a list of future research recommendations.

Chapter 2

Atmospheric Data

2.1 Introduction

This chapter explains contrail formation and persistence, and presents sources of atmospheric data used for contrail and thunderstorm modeling. Contrails are line-shaped clouds that are produced by aircraft engine exhaust at high altitudes. Once a contrail is formed, the evolution of the contrail depends on the ambient temperature and amount of moisture in the air. Thunderstorms are weather events associated with very unstable air, high winds, and usually heavy precipitation. In general, they are caused when a moist, unstable atmosphere exists and there is a rising motion, or lifting force in the atmosphere.

2.2 Contrail Formation

According to the EPA Aircraft Contrails Factsheet, contrails are line-shaped clouds that typically form in the upper atmosphere behind jet aircraft engines and are composed of ice particles [9]. The lifespan of contrails is generally short when they are formed in dry air [9], but contrails persist if they are formed in air that is saturated with respect to ice. The criteria used to determine contrail formation were developed independently by Schmidt and Appleman in 1953, and are now known as the Schmidt-Appleman procedure [51]. The procedure states that a contrail forms when the exhaust and entrained air pass through a thermodynamic state that is first saturated with respect to water, and then into a state that is saturated with respect to ice. The contrail persists if the ambient air is saturated with respect to ice. Figure 2.1 illustrates the Schmidt-Appleman procedure for the case of a persistent contrail [21]. The red dot shows the ambient conditions and the contrail factor line represents the states the exhaust reaches as it evolves from the engine to the ambient air. In the example of Fig 2.1, the red line enters the area of water saturation (contrail forms) and terminates with ambient conditions in the area of ice saturation (contrail persists). It is interesting to note that the slope of the contrail factor line is dependent on the fuel efficiency of the engine which creates the exhaust. Schumann hypothesized, and showed that the more efficient the engine, the steeper the slope of the contrail factor line [52]. This means that the more modern and fuel efficient engines are more prone to creating persistent contrails than engines using older technology. Figure 2.2 shows that a steeper contrail factor line corresponds to a greater likelihood that a contrail is produced [21].

An area in the atmosphere where the air is saturated with respect to ice is called an ice supersaturated region (ISSR) and the relative humidity with respect to ice (RHi) in this region is greater than or equal to 100%. In general, ISSRs are have a very large lateral extent, but are relatively thin vertically. It has been observed that regions of ice supersaturation are defined by boundaries on the order of 150 km horizontally and approximately 500 m vertically [53,54]. Figure 2.3 shows the RHi field over the continental US for a specific day at four times. It has been observed, and can be seen in Fig. 2.3 that fields of RHi have slow dynamics compared to other atmosphere phenomena such as thunderstorms. Additionally, the likelihood that contrails form depends on the time of day as well as the time of year. Stuber et al. found that the formation of contrails can be correlated to the diurnal and annual cycle of air traffic [55]. They found that while night flights account for only 25% of daily flights, they contribute 60%-80% of the contrail radiative forcing. Stuber et al. also observed that while winter flights only account for 22% of annual air traffic, they contribute 50% of the annual contrail formation [55].

Over time, line shaped contrails either dissipate or evolve into a layer of cirrus clouds known as contrail induced cirrus [12]. The level of understanding of the physics and effect of contrail induced cirrus on climate is low, but it is hypothesized that contrail induced cirrus are the dominating impact from contrails [56]. For the purposes of this dissertation it was assumed that contrail induced cirrus scales linearly with the number of persistent contrails produced, and that their evolution is independent of time, place, and the atmospheric conditions surrounding the persistent contrail.

2.3 Relative Humidity Data

In order to model the formation of persistent contrails, the field of relative humidity with respect to ice needs to be known. Currently, areas of ice supersaturation can be extracted from global atmosphere models, but these data are not very accurate. Additionally, it is difficult to measure relative humidity at altitudes near cruise due to the extremely cold temperatures [12]. This section presents sources of RHi data, and the global forecast models that are used in this dissertation.

2.3.1 Sources of Data

Tools such as the Microwave Limb Sounder (MLS) and the Atmospheric Infrared Sounder (AIRS) have been used to detect areas of ice super saturation [57, 58]. However, the data available from these tools lack adequate vertical resolution and depth to describe the vertical profile of the supersaturation in the region. The MOZAIC program collected 9 years of data along major flight routes [59]. Additionally, satellites have been used to detect contrail formation but not explicitly areas of supersaturation. The European Center for Medium-Range Weather Forecasts (ECMWF) provides a global model for relative humidity estimates [60]. The Rapid Update Cycle (RUC) is an atmospheric prediction system that is principally a numerical forecast model developed for users needing short-range weather forecasts [61]. The RUC data have horizontal resolutions of 20km, 40km, and 60km. The vertical resolution of the data is isobaric pressure levels ranging from 100-1000mb in 25mb increments.

The sources of data for relative humidity at high altitudes are, in general, not suited to represent ISSRs for the purpose of path planning. The ideal data set for this research would be a three-dimensional grid of fine vertical resolution. The MLS, AIRS, and MOZAIC data are too sparse to fill the grid required by the path planner. The ECMWF and RUC forecasts provide an adequate grid of data, but the accuracy of the data is questionable. The next section will discuss how the research in this dissertation dealt with these issues, and describe the assumptions that were made.

2.3.2 Relative Humidity Model

As stated previously, an accurate relative humidity forecast model does not exist to predict areas of ice supersaturation. Additionally, the available RHi datasets are too sparse to create the grid of RHi values needed to complete this research. However, ISSRs are observed to have distinct characteristics such as a large lateral area and thin vertical profile. Therefore, if a representative dataset at least emulates these characteristics, it can be used to develop a path planner for persistent contrail avoidance. The assumption is that once an accurate forecast model becomes available, it can replace the representative dataset, and in turn increase the accuracy of the path planner. It was found that Rapid Update Cycle relative humidity data share similar characteristics as observed fields of RHi, and were therefore were selected as the representative dataset. This dataset, albeit inaccurate, provides the necessary foundation for the development of the path planning algorithm developed in this dissertation.

The RUC data have horizontal resolutions of 20km, 40km, and 60km, and a vertical resolution of 25mb. The RUC does not directly output RHi, but it can be calculated from the relative humidity with respect to water (RHw) and the environmental temperature. This dissertation primarily used archived RUC data from November 17, 2001.

Figure 2.3 shows the November 17, 2001 field of RHi at different times. The red areas indicate RHi greater than 100% in the region. During the time period of 1700Z-2000Z the RHi field did not change significantly, which is typical. Also, it should be noted that the size of the RHi fields change with altitude. For MILP implementation, the areas of RHi > 100% were represented as overlapping cuboids. Figure 2.4 shows a RHi field and its three dimensional representation as a set of cuboids.

It should be mentioned again that while the RHi estimates from the RUC data are representative of areas of supersaturation, the values of RHi in the estimates are biased. This is a result of dry bias in some radiosonde measurements that are used in the model [12]. The assumption of this dissertation is that even though the RHi fields might not be accurate, they are representative enough that a path planning tool can be developed and successfully applied to a more accurate dataset when it comes available.

2.4 Convective Weather Data

Thunderstorms are considered a hazard to flight due to the severe wind shear and turbulence associated with their penetration. The evolution and movement of a thunderstorm is still a very active research topic and there exists significant uncertainty in forecasted data. The largest uncertainty in convective forecasts is so called “area probability”, which is the probability that a thunderstorm will occur in a specific area at a specific time [62]. This uncertainty can cause unnecessary delays in the NAS due to conservative decision making by air traffic managers [62]. An on-board trajectory planner would ease the impact of this uncertainty, and also provide more efficient routes around the storms.

2.4.1 Sources of Data

There are generally two types of convective weather data: radar observations and thunderstorm forecasts. Radar observations are the most accurate form of data, but radar data alone can not be used to predict the movement and evolution of a thunderstorm. Thunderstorm forecasts provide a prediction of the movement of a thunderstorm, but are not updated as frequently as radar data.

The Collaborative Convective Forecast Product is a manually generated 2, 4, and 6 hour forecast issued every 2 hours and distributed by the National Weather Service (NWS) Aviation Weather Center [63]. This forecast provides a graphical representation of areas of potential convection, however, it is susceptible to forecasting errors and difficulty of translating the forecast information into something usable by human air traffic managers for airspace capacity considerations. The National Convective Weather Forecast (NCWF), the Regional Convective Weather Forecast (RCWF), and the Terminal Convective Weather Forecast (TCWF) models are automated forecast products used at various Air Traffic Control (ATC) facilities [63].

The National Center for Atmospheric Research (NCAR) is responsible for a database of radar data called National Convective Weather Detection. These data are obtained using the Vertically Integrated Liquid (VIL) algorithm, which converts radar reflectivity to a measure of liquid water content. VIL can be correlated to a Video Integrator and Processor (VIP) rating, which is a discretization of storm strength from 0-6 [64]. This is commonly seen on radar images as the color coding of thunderstorm strength.

2.4.2 Convective Weather Model

The convective weather model developed for this research was designed to provide a frequently updated radar picture and a simple model of thunderstorm evolution. Radar data were acquired from NCAR, and the data were converted to VIP ratings representing the storm strength. It has been observed that pilots avoid thunderstorm cells with a VIP rating of 3 or greater, so these regions were represented by cuboids, and flight was precluded through these areas. Figure 2.5 is an example of the convective weather data used in this dissertation.

Storm movement and evolution were predicted with simple model based on the size of the thunderstorm cell and its observed movement in the past time step. This was designed to be very simple procedure to ease the computational burden of the on-board path planner. The general procedure can be described as:

1. Enlarge the avoidance area.
2. Find centroid of each avoidance area.
3. Project the centroid of each area forward in time based in the movement history in the previous two time steps.

The avoidance area is enlarged to add a safety margin to account for uncertainty in the forecast model. The amount the region is enlarged depends on the amount of risk allowed by the flight, and this is a user defined input to the path planner. For the cuboid representation used in the dissertation, the centroid is found with Equation 2.1

$$\left(x_c, y_c, z_c\right) = \left(\frac{x_{high} - x_{low}}{2}, \frac{y_{high} - y_{low}}{2}, \frac{z_{high} - z_{low}}{2}\right) \tag{2.1}$$

where (x_c, y_c, z_c) is the location of the centroid, and the high and low subscripts refer to the maximum and minimum values of the cuboid in the x, y, and z directions. The velocity of the centroid was found with Equation 2.2

$$\left(v_{x,c}, v_{y,c}, v_{z,c}\right) = \left(\frac{x_{c,i} - x_{c,i-1}}{dt}, \frac{y_{c,i} - y_{c,i-1}}{dt}, \frac{z_{c,i} - z_{c,i-1}}{dt}\right) \tag{2.2}$$

where ($v_{x,c}$, $v_{y,c}$, $v_{z,c}$) is the velocity of the centroid, and *dt* is the time step. The future position of the thunderstorm is simply calculated with Equation 2.3

$$\left(x_{c,i}, y_{c,i}, z_{c,i}\right) = \left(x_{c,0} + i \cdot v_{x,c} dt, y_{c,0} + i \cdot v_{y,c} dt, z_{c,0} + i \cdot v_{z,c} dt\right) \tag{2.3}$$

where ($x_{c,i}$, $y_{c,i}$, $z_{c,i}$) is the position of the centroid at time step *i*, and ($x_{c,0}$, $y_{c,0}$, $z_{c,0}$) is the present position centroid.

2.5 Summary

This chapter presented the atmospheric data used in the path planning algorithm developed by this research. The reported mechanism for contrail formation was explained, and sources of relative humidity data were discussed. Rapid Update Cycle data were selected to model relative humidity because the data represented the general characteristics of a relative humidity field, albeit inaccurately. An assumption inherent in this dissertation allows for the fact that even though the current models of RHi fields may not be accurate, they are representative enough that a path planning tool can be developed, and successfully applied to more accurate data when it comes available.

Convective weather was discussed in terms of both radar reports and radar forecasts. A simple thunderstorm propagation model was presented to predict the evolution of a storm with a specified margin of safety. The simplicity of this model is required due to the on-board nature of the path planning algorithm, which requires computation to adhere to real time constraints.

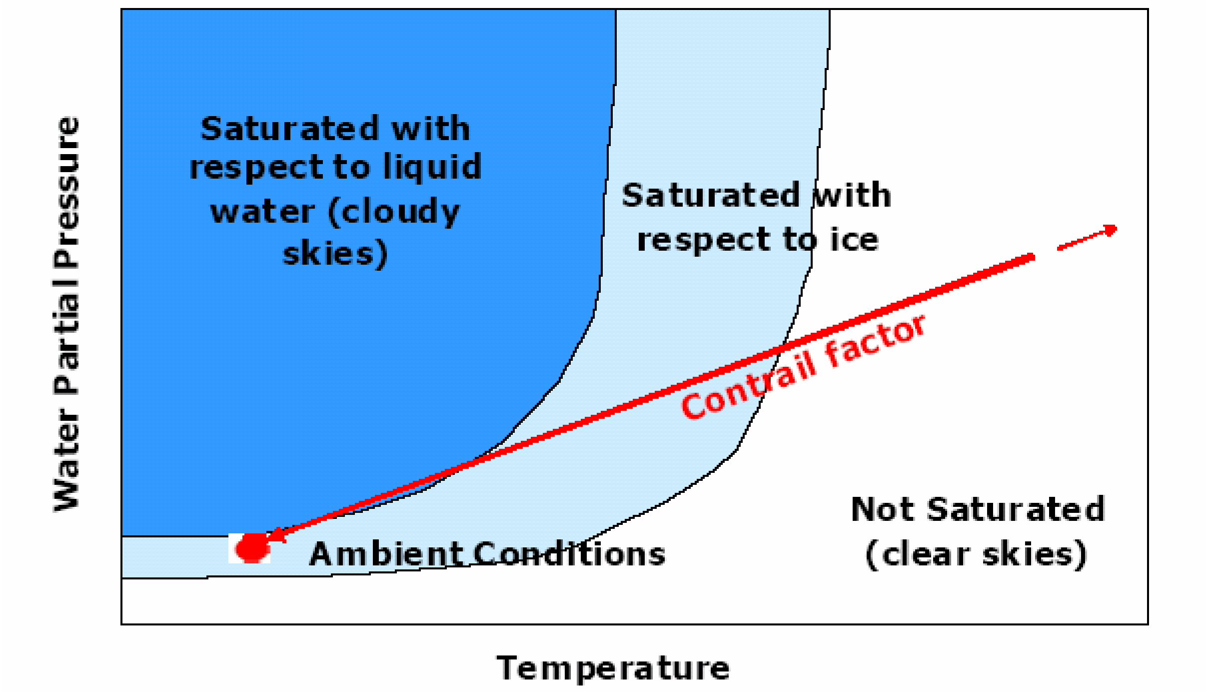

Figure 2.1. Illustration of the Schmidt-Appleman procedure [21].

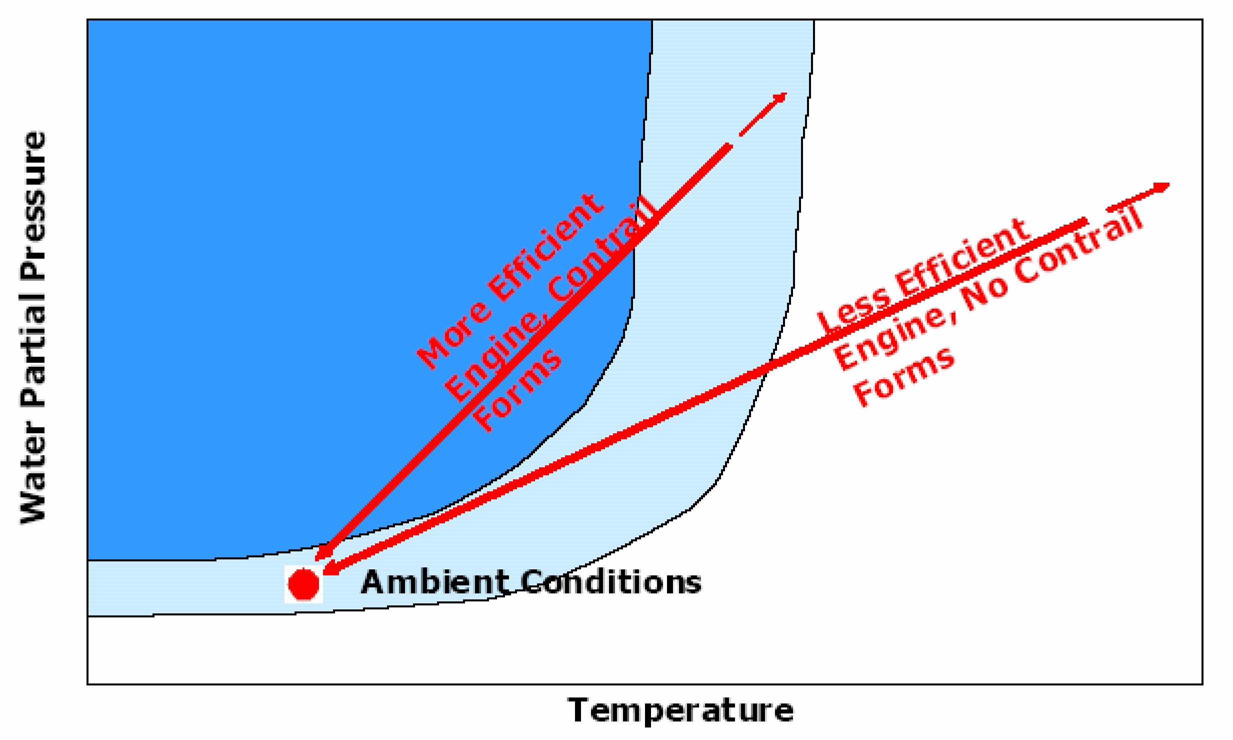

Figure 2.2. Illustration of the Schumann hypothesis [21].

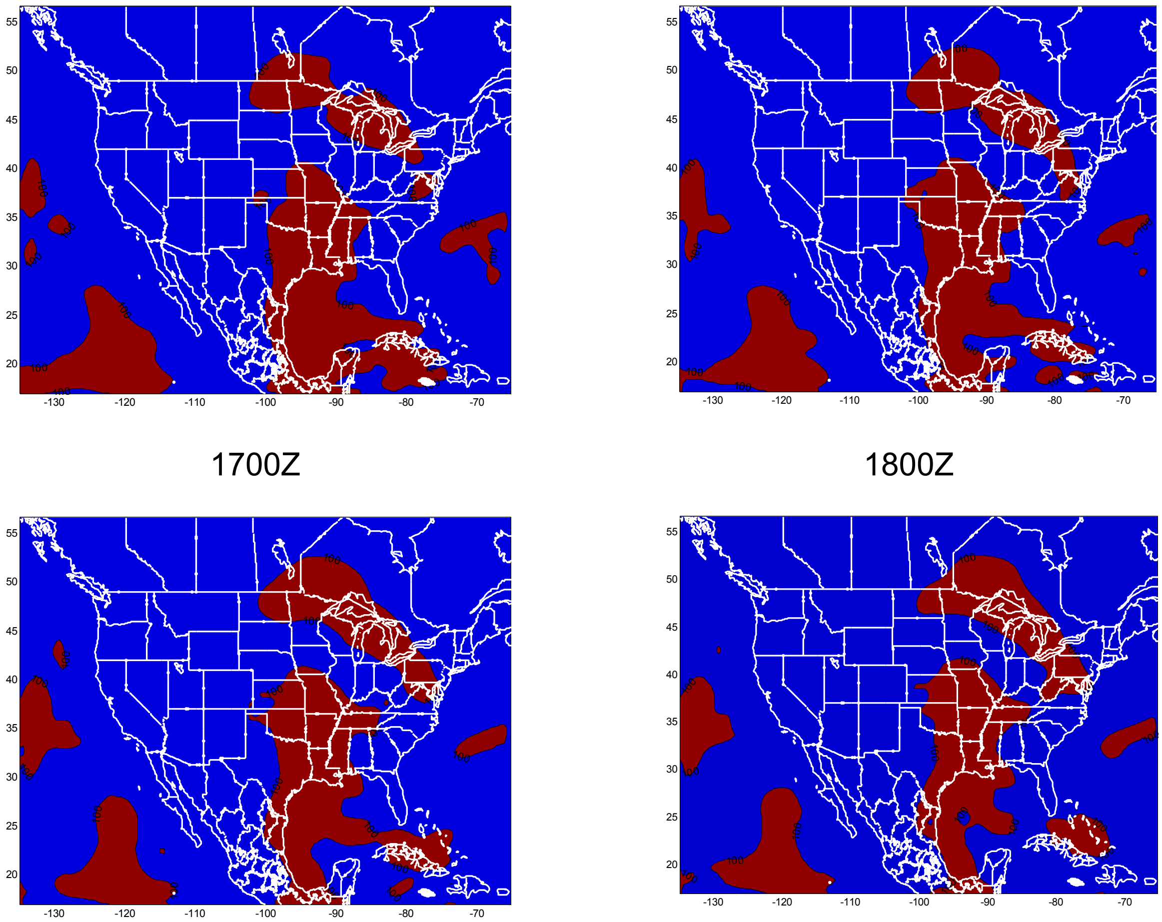

1700Z 1800Z

1900Z 2000Z

Figure 2.3 Fields of RHi at Different Times on November 17, 2001, RUC data.

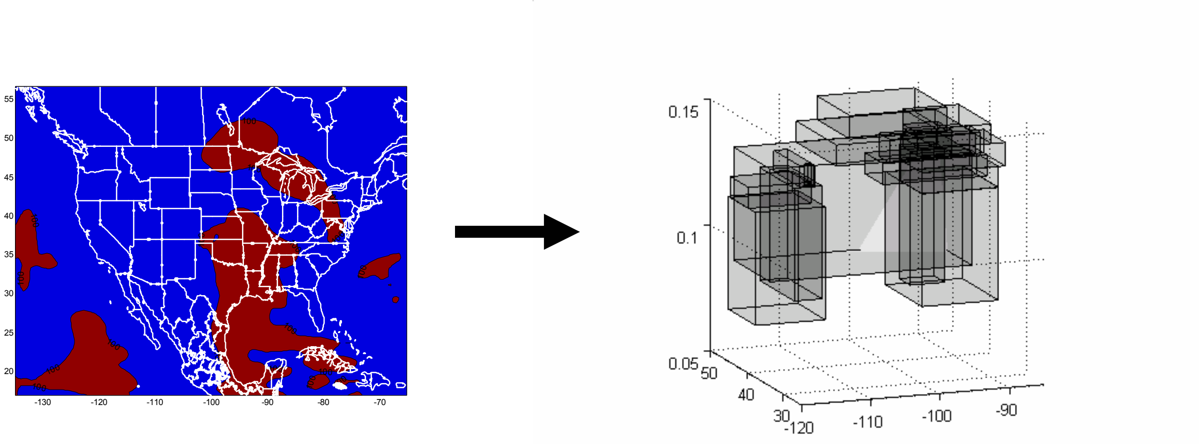

Figure 2.4. Cuboid representation of RHi data for implementation into the MILP.

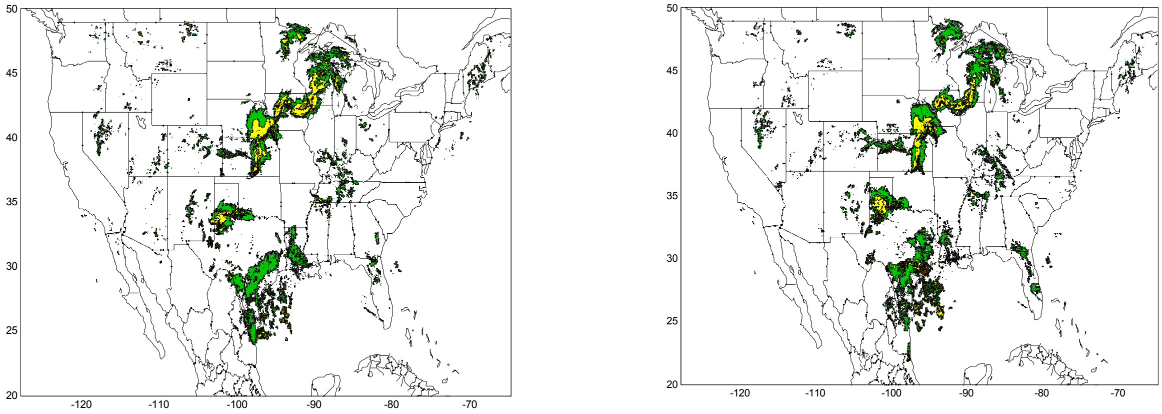

1700Z 1800Z

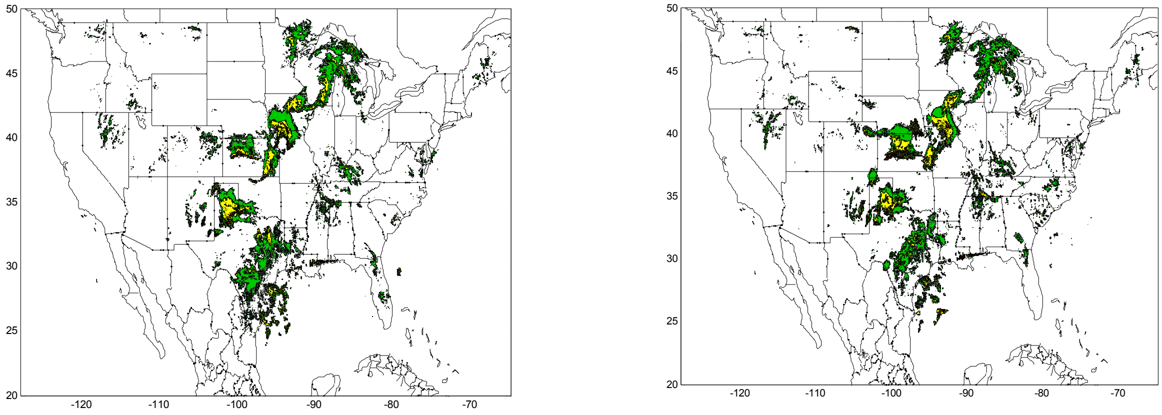

1900Z 2000Z

Figure 2.5. Convective Weather Patterns at Different Times on May 7, 2008. The yellow areas indicate VIP ≥ 3.

Chapter 3

Three-Dimensional Trajectory Generation

3.1 Introduction

This chapter presents a three-dimensional path planning algorithm based on mixed-integer programming (MIP). This approach is similar to existing trajectory planning algorithms [40, 41] which use receding horizon mixed-integer linear programming (MILP) to generate a trajectory in an environment of obstacles. MIP is well suited to the problem due to its ability to directly include logical constraints for obstacle representation, and because it employs an optimization framework that can easily handle the dynamical and performance constraints used in the aircraft and fuel burn models. To improve the computational performance of the optimization, a receding horizon strategy was employed. The receding horizon controller solves the MILP for a predetermined number of time steps toward the destination and the remaining trajectory is accounted for with a cost-to-go path approximation. The main contribution of the chapter is the application of MIP to a large-scale aircraft path planning problem with a realistic aircraft fuel burn and performance model. The aircraft fuel burn model is based on real aircraft data and engine simulations, and is more advanced than previous models in the MIP framework. Additionally, the difference between linear and quadratic cost functions is explored through simulation.

3.2 Aircraft Modeling

The accuracy of an optimization is directly related to the quality of the models used in the calculations. The objective of this path planning problem is to find minimum fuel trajectories around disturbances in the environment such as persistent contrail formation and thunderstorms. The cost function in this optimization is primarily a function of aircraft fuel burn, which is modeled by the combination of an aircraft performance model and an engine performance model.

Aircraft fuel burn is a complicated quantity that is dependent on many states and aircraft-specific parameters. In practice, fuel burn is calculated based on flight test data taken over a wide range of operating conditions. In lieu of these data, fuel burn can be approximated with a limited set of aircraft and engine data. Previous studies have predicted fuel burn using either the FAA's System for Assessing Global Emissions (SAGE) [21], a modified version of SAGE [65], or a quadratic approximation based on velocity [66]. Existing works using the MILP framework have approximated fuel burn by the 1-norm of the aircraft acceleration [43, 44]. This research uses aircraft data and engine performance software to approximate fuel burn over the cruise flight envelope as a function of altitude, velocity, and acceleration. This method is a more accurate model than the 1-norm of acceleration or quadratic approximation of velocity, but is still simple enough to implement in the mixed-integer programming framework.

3.2.1 Aircraft Performance Model

The aircraft performance model used for this research is intended to emulate the en-route performance characteristics of medium-range aircraft such as the Boeing 737 and

Airbus A320. The following restrictions are placed on the altitude and Mach number to confine aircraft performance to the cruise envelope, as seen in Eq. 3.1

$$\begin{aligned} &0.70 \leq Ma \leq 0.82 \\ &28{,}000\,ft \leq z \leq 42{,}000\,ft \end{aligned} \tag{3.1}$$

where *Ma* is the Mach number and *z* is the altitude.

To compute the performance, the drag coefficient is extracted from drag polar data [67, 68] for a range of Mach numbers, altitudes, and weights. Figure 3.1 depicts the drag polar data for flight conditions within the range of Eq. 3.1. The thrust required is calculated for the range of Mach numbers and altitudes inside the cruise flight envelope, and for three weights, each representing the aircraft weight at a different fuel state along the flight path. The thrust required for level, steady flight is calculated with Eq. 3.2

$$T_{req} = q \cdot C_D \cdot S \tag{3.2}$$

where *q* is dynamic pressure, C_D is the drag coefficient, and *S* is the planform area. Figure 3.2 shows how the thrust required changes with altitude and Mach number for a given weight.

During cruise, commercial aircraft frequently change altitude using the flight level change mode of the flight management system, and therefore it is assumed that the thrust is set to *maximum climb thrust* during climb, and *idle* during descent. The climb performance is limited by a maximum rate of climb constraint. The maximum rate of climb (R/C_{max}) is determined with Eq. 3.4

$$R/C_{\max} = \frac{V_\infty \left(T_A - T_{req}\right)}{W} \tag{3.4}$$

where V_∞ is the velocity of the aircraft, T_A is the thrust available, T_{req} is the thrust required in level flight at this flight condition, and W is the weight of the aircraft. The thrust available at altitude is approximated by the Eq. 3.5 [68]

$$T_A = (T_A)_0 \cdot \left[\frac{\rho}{\rho_0}\right]^m \tag{3.5}$$

where $(T_A)_0$ is the thrust available at sea level, ρ_0 is standard sea-level density, and m is a thrust factor. Figure 3.3 shows the maximum rate of climb generated by the performance model. The thrust available at sea-level was obtained from the engine model, which is presented in the next section, and the thrust factor was approximated to be 0.8. The absolute ceiling for each weight is equal to the x-intercept of the lines in the figure.

3.2.2 Engine Performance Model

Engine performance is obtained with the Engine Performance Analysis Program v4.2 [69]. This program uses a set of engine parameter inputs and flight conditions (Mach number, altitude) to compute curves of Thrust Specific Fuel Consumption (TSFC) vs. engine thrust. The engine is assumed to be a high-bypass turbofan, and the software input parameters are given in Table 3.1. The program was run for the range of Mach numbers and altitudes given by Eq. 3.1. The engine performance data are tabulated for use in the aircraft fuel burn model.

Figure 3.4 shows the relationship between TSFC and altitude for different Mach numbers in steady, level flight. The TSFC trends are as expected with lower TSFCs corresponding to the higher altitudes and lower Mach numbers. The outputs of the engine performance model were combined with the outputs of the aircraft performance model to approximate the relationship between fuel burn, Mach number, and altitude.

3.2.3 Fuel Burn Model

The fuel burn model developed by this research is a combination of the aircraft performance and engine performance models presented previously. Aircraft fuel burn is calculated for a range of Mach numbers, altitudes, and weights using Eq. 3.6

$$W_f = T_{req} \cdot TSFC \tag{3.6}$$

where W_f is the fuel flow, T_{req} is the thrust required, and *TSFC* is found from the engine performance model. Figures 3.5 and 3.6 show the relationship of altitude and velocity with fuel flow, respectively, for an aircraft weight of 145,000 lb with engine and drag characteristics described above.

Figures 3.5 and 3.6 show the highly nonlinear nature of fuel burn. To simplify this behavior, a nominal flight condition is selected using Mach number and altitude for a given weight, and the change in fuel burn around the nominal flight condition is modeled. The nominal cruise Mach number is selected to be 0.78 based on typical cruise speeds for the type of aircraft considered in this model [70]. The nominal cruise altitude is selected based on the fuel-optimal altitude for a Mach number of 0.78 and for a given weight. For a true minimum fuel solution, the nominal altitude should increase continuously as fuel is burned, resulting in a cruise climb flight profile. However, this procedure is generally not performed in practice because of air traffic control restrictions. Instead, a step-climb procedure is used, where the altitude is increased in discrete steps along the flight path [71]. This fuel burn model emulates a step climb by using three nominal cruise altitudes based on the optimal altitude for three aircraft weights. The initial weight is assumed to be 145,000 lbs, and the subsequent weights are 135,000 lbs and 125,000 lbs, which correspond to optimal altitudes of approximately 34,000 ft, 36,000 ft, and 38,000 ft, respectively. Table 3.2 gives the nominal flight conditions used to develop the model.

Figures 3.7 and 3.8 show the sensitivity of fuel flow to changes in altitude and velocity around the nominal flight conditions, respectively. In Fig. 3.7, the general trend of all three curves is that fuel burn increases for flight at altitudes both above and below the nominal altitude. At altitudes below the nominal altitude, the fuel burn approximation in Eq. 3.6 is driven higher by increasing TSFC. At altitudes above the nominal altitude, higher thrust is required to maintain level flight, and in turn increases fuel burn. The trends exhibited in Fig. 3.8 are as expected, where higher velocity corresponds to higher fuel consumption. These curves were converted to piecewise linear and quadratic functions for implementation into the path planning algorithm. The cost function is described in more detail in Section 3.3.3, and a quantitative comparison is given in Section 3.4.

It should be noted that this aircraft fuel burn model is developed assuming level, steady flight in the nominal flight conditions. This is a valid assumption considering the highly unaccelerated nature of en-route commercial operations. Accelerated flight is modeled in a first order fashion, and is described in more detail in Chapter 3.3.4.1. Climbs and descents were modeled with the assumption that commercial aircraft frequently change altitude using the flight level change mode of the flight management system, and therefore it was assumed that the thrust is set to *maximum climb thrust* during climb, and *idle* during descent.

3.3 Path Optimization

Recently, the field of path optimization has blossomed due to the increased use of Unmanned Aerial Vehicles (UAVs) and the desire for more autonomy in the National Airspace System (NAS). Existing path planning research has been primarily concerned with the two dimensional motion of an aircraft or ground vehicle subject to relatively simple cost and performance constraints. This section describes the formulation of a three-dimensional path planning algorithm that is subject to a realistic aircraft model and capable to be solved on-board an aircraft in real time.

3.3.1 Mixed-Integer Programming

The field of path planning is broad and many techniques have been studied in the literature [29-45]. Of these techniques, mixed-integer programming (MIP) shows the most promise as a technique for the objective of this research. This is because MIP can directly incorporate logical constraints such as obstacle avoidance, and provide an optimization framework that employs dynamic constraints such as turn and maximum rate of climb limitations [40, 41]. The MIP framework allows both linear and quadratic cost functions, as long as the quadratic term is positive semi-definite. A MIP optimization with a linear cost is called a mixed-integer linear programming (MILP) problem, and with a quadratic cost is called mixed-integer quadratic programming (MIQP) problem. The basic format of a MILP optimization is given as Eq. 3.7

$$\begin{aligned} &\text{minimize} : c^T x \\ &\text{subject to} : Ax \le b \\ &\quad l \le x_1 \le u \\ &\quad x_2 \text{ binary} \end{aligned} \tag{3.7}$$

where x is composed of x_1 and x_2, and x_1 is a continuous variable bounded by l and u, and x_2 is a binary variable. The vector c describes the cost, and the matrix A and vector b express the system constraints. The basic form of a MIQP is given as Eq. 3.8

$$\begin{aligned} &\text{minimize} : \frac{1}{2} x^T Q x + d^T x \\ &\text{subject to} : Ax \le b \\ &\quad l \le x_1 \le u \\ &\quad x_2 \text{ binary} \end{aligned} \tag{3.8}$$

where x is composed of x_1 and x_2, and x_1 is a continuous variable bounded by l and u, and x_2 is a binary variable. The matrix Q describes the cost and must be symmetric and positive semi-definite. The matrix A and the vector b express the system constraints.

The MILP problem has been proven in the literature to be NP-Hard in the number of binary variables [46, 47], which informally means that the problem is at least as hard as one that is NP-Complete, where the solutions to NP-Complete problems can be verified in polynomial time. The significance of this is that as the number of binary variables in the problem grow, the computation time increases dramatically. To mitigate the computational burden imposed by increasing problem size, a receding horizon approach is used to break the MIP into a detailed trajectory optimization close to the aircraft and a coarse cost-to-go approximation to the destination.

3.3.2 Receding Horizon Control

The basic approach of receding horizon control is to utilize a model of the system to forecast future behavior and to make control inputs that optimize the performance objectives and satisfy the constraints. The optimization is repeated online as new information about the aircraft states and the environment is obtained. Each online optimization uses information on the current environment, which is continually updated based on the latest observations of the aircraft.

Receding horizon control is also known as model predictive control (MPC), and has been successfully applied to the field of process control [72, 73] where the dynamics of the system are typically on the order of minutes or hours. Recent advances in computational power have enabled the application of receding horizon control to systems with faster dynamics, such as cars and aircraft [74].

The primary purpose of receding horizon control in this dissertation is to reduce the computational burden of the optimization. To do this, the receding horizon controller truncates the optimization at a finite horizon and uses a terminal penalty, also known as cost-to-go, to represent the remaining trajectory. This is a technique of multi-resolution planning, where the path planner optimizes in detail over the local horizon and connects the detailed trajectory to a more coarse approximation to the goal. The trajectory is broken into three distinct segments: the detailed trajectory, the line of sight, and the cost-to-go. The detailed trajectory is a high resolution optimization of the trajectory subject to the realistic aircraft model presented earlier, and also called the planning horizon. The cost-to-go is a coarse approximation of the trajectory through a cost grid, and the line of sight connects the end of the detailed trajectory to the start of the cost-to-go. The Figure 3.9 shows the segments of the receding horizon trajectory in detail, with the planning horizon represented by the green line, the line-of-sight by the blue line, and the cost-to-go by the magenta line.

The receding horizon controller solves the MIP for a predetermined number of time steps toward the destination and the procedure is repeated, with the remaining trajectory accounted for with the cost-to-go, until the destination is reached. Compared to non-receding horizon MIP, which generates a full detailed trajectory up to the destination, this approach does not waste computational resources on a detailed plan in the far future, where not much information is available and large changes in the environment are possible.

3.3.3 Constraints and Limitations

This section describes the construction of the constraints and limitations imposed in the MIP optimizations. Specifically, it discusses the implementation of the atmospheric model introduced in Chapter 2 and the aircraft fuel burn and performance model established in Chapter 3.2. The order of the section is as follows: first, the dynamical and aircraft performance constraints are discussed in detail. Then, the obstacle avoidance constraints are explored and the section concludes with a detailed explanation of the implementation of the cost function.

3.3.3.1 Dynamical and Aircraft Performance

The dynamics of aircraft are governed by coupled non-linear differential equations, and are therefore very difficult to implement in an on-board path planner. To simplify the equations of motion, the aircraft is treated as a simple point mass subject to realistic aircraft performance limits and constraints. While this method is not the most accurate, it provides a good representation of the flight envelope, and occurs commonly in path planning [75].

The dynamical constraints presented here are that of a double integrator and the evolution of the aircraft states is governed by Eq. 3.9

$$\mathbf{x}(k+1) = A\mathbf{x}(k) + b\mathbf{u}(k) \tag{3.9}$$

where,

$$\mathbf{x} = \begin{bmatrix} x \\ y \\ z \\ v_x \\ v_y \\ v_z \end{bmatrix} \qquad \mathbf{u} = \begin{bmatrix} a_x \\ a_y \\ a_z \end{bmatrix} \qquad A = \begin{bmatrix} I_3 & dt \cdot I_3 \\ 0_3 & I_3 \end{bmatrix} \qquad b = \begin{bmatrix} \tfrac{1}{2} dt^2 \cdot I_3 \\ dt \cdot I_3 \end{bmatrix} \tag{3.10}$$

The vector **x** represents the position and velocity of the aircraft, the vector **u** represents the acceleration, *k* is the discrete time index, and *dt* is the size of the time step. The matrix I_3 represents a 3x3 identity matrix, and **u** is taken along the (x,y,z) basis.

The calculation of the magnitude of the velocity from the individual velocity components requires the square root of the sum of the squares, which is obviously a nonlinear operation and inadmissible in MILP. Therefore an accurate method is needed to approximate the magnitude of velocity because fuel burn is largely dependent on the speed of the aircraft. The following procedure provides a linear approximation of the velocity and acceleration for implementation into the MILP [48]

$$V[k] \geq v_x[k]\cos\left(\frac{2\pi w}{M}\right) + v_y[k]\sin\left(\frac{2\pi w}{M}\right) \tag{3.11}$$

$$V[k] - R \cdot (1 - b_v[k,w]) \leq \left(v_x[k]\cos\left(\frac{2\pi w}{M}\right) + v_y[k]\sin\left(\frac{2\pi w}{M}\right)\right) \cdot 1.01 \tag{3.12}$$

$$A[k] \geq a_x[k]\cos\left(\frac{2\pi w}{M}\right) + a_y[k]\sin\left(\frac{2\pi w}{M}\right) \tag{3.13}$$

where V and A are the velocity and acceleration magnitudes respectively, M is the number of faces of the polygon, w is an integer in the set of integers from 1 to M, $b_v[k,w]$ are binary variables, and R is a large constant. M is a user defined variable where the approximation accuracy increases with M. Equation 3.12 is necessary to include non-convex constraints on minimum velocity. Figure 3.10 shows how the approximation fits inside of the velocity circle. It should be noted here that the velocity and acceleration magnitude approximations only account for motion in the x-y direction. Motion in the z-direction is assumed to be very small compared to motion in the x-y direction, and is treated separately. Additionally, V and A are bounded to constrain the optimization within a realistic flight envelope, as given in Eqs. 3.13 and 3.14

$$V_{\min} \leq V \leq V_{\max} \tag{3.13}$$

$$A \leq A_{\max} \tag{3.14}$$

The vertical velocity is constrained by a rate of climb limit, which is a function of altitude. The maximum rate of climb is defined by Eq. 3.15

$$RC_{\max} + \beta_1 z = \beta_0 \tag{3.15}$$

where RC_{max} is the maximum rate of climb, and β_i is computed in the aircraft model. The rate of climb is constrained by the maximum rate of climb using Eq. 3.16

$$v_z - RC_{\max} \leq 0 \tag{3.16}$$

The maximum descent rate is represented by a lower bound on the vertical velocity and is chosen to be a value consistent with normal operation of commercial aircraft [68].

3.3.3.2 Avoidance Constraints

As discussed in Chapter 2, persistent contrails form in specific areas of the atmosphere where the relative humidity with respect to ice is greater than 100%. Additionally, aircraft are forced to divert around thunderstorms that are above a certain strength and altitude. These disturbances to the National Airspace System are represented as cuboids in the optimization, and flight through these cuboids is precluded. One benefit of the MIP framework is that it can handle hard avoidance constraints, which strictly enforce that all feasible trajectories remain outside of the cuboid.

Persistent contrail, and thunderstorm avoidance is accomplished using hard constraints of the following form:

$$\begin{aligned} x[k] &\leq B_{x,low} + R \cdot b_{o,1}[k] \\ x[k] &\geq B_{x,high} - R \cdot b_{o,2}[k] \\ y[k] &\leq B_{y,low} + R \cdot b_{o,3}[k] \\ y[k] &\geq B_{y,high} - R \cdot b_{o,4}[k] \\ z[k] &\leq B_{z,low} + R \cdot b_{o,5}[k] \\ z[k] &\geq B_{z,high} - R \cdot b_{o,6}[k] \end{aligned} \tag{3.17}$$

$$\sum_{i=1}^{6} b_{o,i}[k] \leq 5 \tag{3.18}$$

where *B* defines the 6 planes that compose the cuboid which represents an area of RHi > 100%. The variable $b_{o,i}$ is a binary variable that is either 1 or 0 depending on whether constraint *i* is active. Eq. 3.18 ensures that all 6 constraints are never all active at the same time, meaning that the aircraft stays out of the area of RHi ≥ 100%, or the thunderstorm cell.

3.3.4 Cost Function

As shown in Chapter 3.2, aircraft fuel burn is extremely nonlinear by nature, and in turn, difficult to implement in a linear structure. In a cruise flight condition, the main drivers of fuel burn are velocity, altitude, and weight. If climbing or descending, the throttle setting is also a factor. To reiterate the assumptions used in Section 3.2, the effect of Mach number, altitude, and throttle setting were assumed to be decoupled around a nominal flight condition and weight. Also, the effect of weight was accounted for by changing the nominal flight condition at pre-specified intervals of time. The time intervals were selected to correspond to the approximate time the aircraft weight reaches 135,000 lbs and 125,000 lbs during the flight. This section will discuss the cost function of the detailed trajectory optimization and then the cost-to-go.

3.3.4.1 Detailed Trajectory

The fuel burn curves given in Figs. 3.7 and 3.8 were written as a set of piecewise linear functions. The piecewise linear representation of fuel burn as a function of weight is described by Eqs. 3.19 and 3.20

$$f_a(W) = \max(\mu_1 z + \mu_2, \ldots, \mu_i z + \mu_{i+1}) \quad \text{for } i = 1,3,5,\ldots \tag{3.19}$$

$$f_v(W) = \max(\sigma_1 V + \sigma_2, \ldots, \sigma_i V + \sigma_{i+1}) \quad \text{for } i = 1,3,5,\ldots \tag{3.20}$$

where *$f_a(W)$* and *$f_v(W)$* are the fuel burn associated with altitude and velocity, respectively, *W* is weight, and μ_i and σ_i determine the piecewise linear function. These equations are written in MILP format as follows:

$$\begin{array}{c} \mu_1 z - f_a \le \mu_2 \\ \vdots \\ \mu_i z - f_a \le \mu_{i+1} \end{array} \quad \text{for } i = 1,3,5,\ldots \tag{3.21}$$

$$\begin{array}{c} \sigma_1 V - f_v \le \sigma_2 \\ \vdots \\ \sigma_i V - f_v \le \sigma_{i+1} \end{array} \quad \text{for } i = 1,3,5,\ldots \tag{3.22}$$

It should be noted that the coefficients *μ* and *σ* depend on the aircraft weight; therefore these equations are updated whenever the aircraft weight is updated.

Another type of fuel burn approximation is a quadratic function. The quadratic representation of fuel burn as a function of weight is given by Eqs. 3.23 and 3.24

$$f_{Q,a}(W) = \eta_1 z^2 - \eta_2 z + \eta_3 \tag{3.23}$$

$$f_{Q,v}(W) = \eta_4 V^2 - \eta_5 V + \eta_6 \tag{3.24}$$

where $f_{Q,a}(W)$ and $f_{Q,v}(W)$ are the altitude and velocity costs, respectively, and η_i is taken from the aircraft model. A quadratic cost function better represents the data; however, the optimization must be solved with MIQP, which is more computationally demanding than MILP. A study showing the results of a tradeoff between the performance and computation times of different types of cost functions is given in Chapter 3.4.

The effect of climb and descent on fuel burn was assumed to be decoupled from the effects of velocity, altitude, and weight. During cruise, commercial aircraft frequently change altitude using the flight level change mode of the flight management system, and therefore it was assumed that the thrust is set to *maximum climb thrust* during climb, and *idle* during descent. The climb and descent state was characterized with Eq. 3.25

$$\begin{aligned} v_z &\leq R \cdot b_{c\lim b} \\ -v_z &\leq R \cdot b_{descend} \end{aligned} \tag{3.25}$$

where v_z is the vertical velocity, R is an arbitrarily large constant, b_{climb} is a binary variable to indicate climb, and $b_{descend}$ is a binary variable to indicate descent.

The last component of the cost is the fuel burn associated with acceleration, which is determined under the assumption that the fuel burn linearly increases with acceleration. In full, the cost function for the detailed trajectory optimization is written as Eq. 3.26

$$J = \sum_{i=1}^{N_p} \left(f_{a,i} + f_{v,i} + A_i + b_{f,i} f_{term,i} + b_{climb,i} f_{descend} + b_{descend,i} f_{descend} \right) \tag{3.26}$$

where f_a and f_v are the fuel cost associated with altitude and velocity, respectively, A is acceleration, f_{climb} is a weighting associated with maximum climb thrust, and $f_{descend}$ is a weighting associated with idle. Again, it should be noted that f_a and f_v change with aircraft weight, which is updated periodically during the receding horizon optimization. The total trajectory cost includes the detailed trajectory cost (Eq. 3.26), and the cost-to-go, which is presented in the next section.

3.3.4.2 Cost-to-Go

As mentioned previously, the detailed trajectory is planned until the end of the planning horizon, and the remaining trajectory to the goal is approximated by the cost-to-go function. The cost-to-go function approximates the fuel required to go from the end of the planning horizon to the goal by creating a cost map containing the fuel to travel from each node in the map to the goal, and an additional fuel cost to connect the detailed trajectory to the cost map.

The cost map, G_{ij}, is a measure of the cost between nodes i and j, and is found using a visibility graph weighted by the distance between nodes and a cost associated with the altitude of the nodes. The nodes in the cost map are made up of the vertices of the obstacles in the environment, in addition to the aircraft position and the destination position.

To assign a realistic cost approximation to the visibility graph, it was weighted by a two part function. The first part approximates the fuel burn between the nodes based on the Euclidian distance between the nodes, and the second part is an altitude penalty based on a quadratic altitude function evaluated at the average altitude of the connecting nodes. Therefore, the cost to travel between nodes *i* and *j* is given by Eqs. 3.27 and 3.28

$$G_{ij} = \lambda \left\| \mathbf{x}_i - \mathbf{x}_j \right\|_2 + \alpha_1 z_{avg}^2 + \alpha_2 z_{avg} + \alpha_3 \tag{3.27}$$

$$z_{avg} = \frac{1}{2}\left(z_i + z_j\right) \tag{3.28}$$

where z_{avg} is the average altitude of the connecting nodes, α_i is determined by the aircraft model, and λ is a constant used to map distance to fuel burn. Dijkstra's Algorithm was applied to the cost map to find the path of minimum fuel cost from each node to the goal [76]. The output is a vector, C_i, which gives the cost to go from each node *i* in the cost map to the goal.

The cost-to-go was completed by connecting the end of the detailed trajectory to a point in the cost map. The position at the end of the detailed trajectory, $\mathbf{x}[N_p]$, and the node in the cost map, $\mathbf{x}_{cp}$, are chosen by the optimization so that they are mutually visible and that the fuel cost required to travel between them is minimized. Visibility is ensured by requiring a set of interpolation points between $\mathbf{x}[N_p]$ and $\mathbf{x}_{cp}$ to remain outside of the obstacle regions. Equation 3.29 is used to select the visible point

$$\sum_{i=1}^{n_{cp}} b_{vis,i}\mathbf{x}_{cp,i} = \mathbf{x}_{vis}$$
$$\sum_{i=1}^{n_{cp}} b_{vis,i} = 1 \tag{3.29}$$

where $b_{vis,i}$ is a binary variable for cost point i, and n_{cp} is the number of nodes. The cost associated with traveling this path is a function of the Euclidian distance and the altitude of the interpolation points. The Euclidian distance, D, was approximated with Eq. 3.30

$$D \geq \lambda\left(\Delta x \sin\left(\frac{2\pi m_1}{M_1}\right)\cos\left(\frac{2\pi m_2}{M_2}\right) + \Delta y \sin\left(\frac{2\pi m_1}{M_1}\right)\sin\left(\frac{2\pi m_2}{M_2}\right) + \Delta z \cos\left(\frac{2\pi m_1}{M_1}\right)\right)$$
$$m_j = 1 \ldots M_j$$
$$j \in [1,2] \tag{3.30}$$

where m_j is the set of integers between 1 and M_j, and Δx, Δy, and Δz are the distances between $\mathbf{x}[\mathrm{N_p}]$ and $\mathbf{x}_{\mathrm{cp}}$ in the x, y, and z direction, respectively. The altitude of the interpolation points was penalized using the same philosophy as in Eq. 3.27, which penalizes flight away from the optimal altitude. The complete cost-to-go function is given in Eq. 3.31

$$J = C^T \cdot b_{vis} + \sum_{i=1}^{n_I} f_{x,i} + D \tag{3.31}$$

where C is the cost vector, $f_{x,i}$ is the interpolation point altitude cost, n_I is the number of interpolation points, and D is the cost of flight between $\mathbf{x}[\mathrm{k+N_p}]$ and $\mathbf{x}_{\mathrm{cp}}$.

3.4 Linear vs. Quadratic Cost

To contrast the differences in using a linear versus a quadratic cost, consider a single flight from O'Hare International Airport (ORD) to Los Angeles International Airport (LAX) using atmospheric data from November 17, 2001. The objective of this example was to find a fuel optimal trajectory using both linear and quadratic cost functions for this route while flying clear of atmospheric areas containing RHi > 100%. The fuel consumption and computational performance of the trajectories are compared to determine the tradeoff between accuracy and running time. The fuel burn cost is derived from the model presented in Chapter 3.2, and the formulation of Chapter 3.3 describes the dynamical and aircraft performance constraints. Table 3.3 lists the receding horizon parameters and the aircraft performance limitations used in this example. The simulations were computed using CPLEX 10.2 [77] on a PC laptop with a 2.16GHz Intel Core2 Duo processor and 2GB of RAM.

Figures 3.11 and 3.12 show trajectories overlaid on two-dimensional contour plots of the RHi field at different altitudes. The blue, magenta, and cyan trajectories correspond to the receding horizon mixed-integer linear programming (RH-MILP) formulations, each of which used a different number of piecewise linear segments in the cost function. The black trajectory was optimized with a RH-MIQP formulation, and the red trajectory represents the non-receding horizon MILP limiting case formulation, where the goal is contained within the planning horizon. The trajectories were initiated at approximately 34,000 ft, and as seen in Fig. 3.11b, a region of RHi > 100% over southern Iowa and Missouri forced the trajectories to adjusts their routes. The piecewise linear trajectories followed an almost identical horizontal flight path to avoid contrail formation. The quadratic trajectory deviated from the piecewise linear trajectories, which was caused by a difference in the sensitivity of altitude change to fuel burn in their respective cost functions. The full horizon trajectory followed a much straighter flight path to LAX due to the fact that it was computed in one step and not in a receding horizon fashion.

Figures 3.13 and 3.14 show the aircraft and computational performance associated with the trajectories presented in Figs. 3.11 and 3.12. Figure 3.13a shows the velocity profiles of the trajectories, and it is easy to see that the velocities corresponding to the receding horizon formulations are very similar. The difference between the RH-MILPs and the MIQP can, at least in part, be explained by the fact that the MIQP is less sensitive to small velocity variations around the nominal velocity. Figure 3.13b shows the altitude profiles of the trajectories, where it can be seen that an area of RHi > 100% was encountered approximately 0.4 hours into the flight. The four and six piece RH-MILP trajectories chose to climb over the area, while the two piece RH-MILP, the RH-MIQP, and the full horizon trajectories chose to descend under the area. The climb seen by all trajectories at 2 hours is due to the preprogrammed step-climb profile discussed earlier. Figure 3.14b shows the computation expense of each method. The two piece RH-MILP was the cheapest, followed by the four and six piece RH-MILPs, and the RH-MIQP was by far the most expensive from a computation standpoint.

Table 3.4 compares the performance of the four trajectories presented by this example. The RH-MIQP showed the best performance in terms of total fuel burn, but was by far the most computationally expensive. The four and six segment RH-MILPs displayed better fuel burn performance than the two segment RH-MILP, and were only slightly more computationally expensive. The receding horizon trajectories burned between 2.88 % and 2.23 % more fuel than the full horizon limiting case. Overall, the full horizon trajectory burned 2.76% more fuel than a trajectory disregarding contrail avoidance. The four-segment piecewise linear approximation showed the best combination of accuracy and computational performance, and is the type of approximation used in the remainder of this dissertation.

3.5 Summary

This chapter presented the three-dimensional path planning architecture used in this dissertation. The aircraft performance and fuel burn models were discussed and the implementation of the models to the MIP was explained. The avoidance constraints were given and discussed. The cost function and cost-to-go were presented, and a trade study of cost function type was explored. The results showed that the 4 segment piecewise linear cost function had the best combination of accuracy and efficiency for the scales of obstacles considered by this dissertation.

Table 3.1. Engine Performance Software Input Parameters

Mass Flow Rate (max throttle, sea level)	779 lbm/sec
Bypass Ratio	5.1
Compressor Pressure Ratio	32.8
Fan Pressure Ratio	2.3

Table 3.2. Nominal Flight Conditions

Weight (lb)	Altitude (ft)	Mach number	True airspeed (knots)
145,000	34,000	0.78	451
135,000	36,000	0.78	447
125,000	38,000	0.78	447

Table 3.3. Receding horizon parameters and aircraft performance limits

Number of steps in planning horizon (N_p)	8
Number of steps in the execution horizon (N_e)	4
Time step size (Δt)	3 min
Maximum en-route velocity (V_{max})	470 knots
Minimum en-route velocity (V_{min})	417 knots
Maximum altitude (z_{max})	42,000 ft
Minimum altitude (z_{min})	28,000 ft

Table 3.4. Comparison of receding horizon trajectory performance

	MILP N_p = 8, N_e = 4 2 Segment	MILP N_p = 8, N_e = 4 4 Segment	MILP N_p = 8, N_e = 4 6 Segment	MIQP N_p = 8, N_e = 4
Max. Velocity	462.3 knots	462.1 knots	462.1 knots	457.3 knots
Avg. Velocity	454.7 knots	454.6 knots	454.6 knots	451.8 knots
Total Fuel Burn	20,824 lbs	20,731 lbs	20,734 lbs	20,695 lbs
% Difference from Full Horizon	2.88 %	2.42 %	2.43 %	2.23 %
Flight Time	3.35 hrs	3.35 hrs	3.35 hrs	3.40 hrs
Max CPU time /step	3.23 sec	8.88 sec	7.27 sec	62.73 sec
Avg CPU time /step	1.94 sec	3.03 sec	2.69 sec	30.83 sec

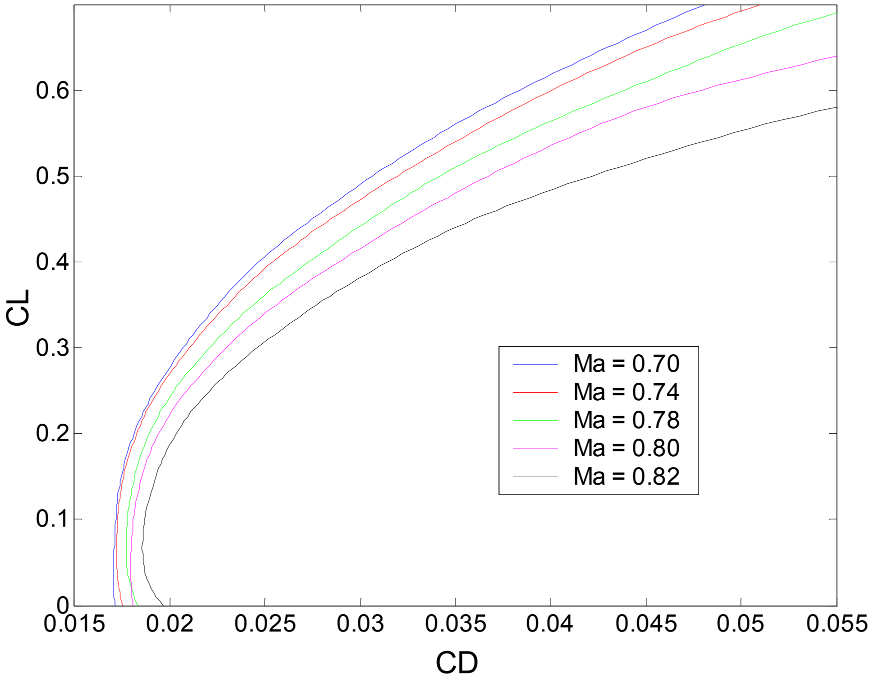

Figure 3.1. Mid-sized jet transport drag polar data for a range of Mach numbers.

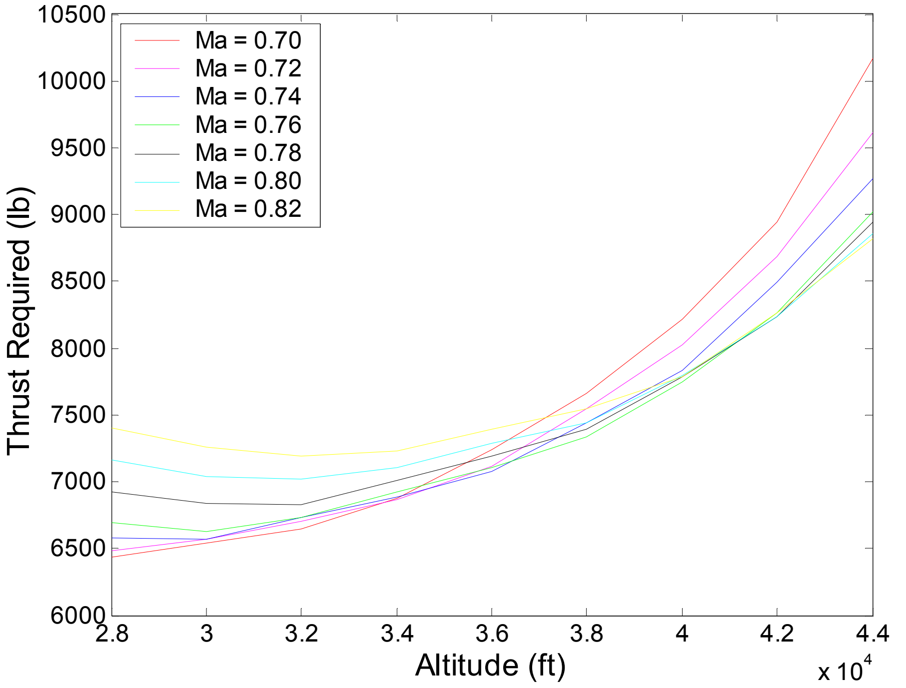

Figure 3.2. Mid-sized jet transport thrust required for a range of Mach numbers.

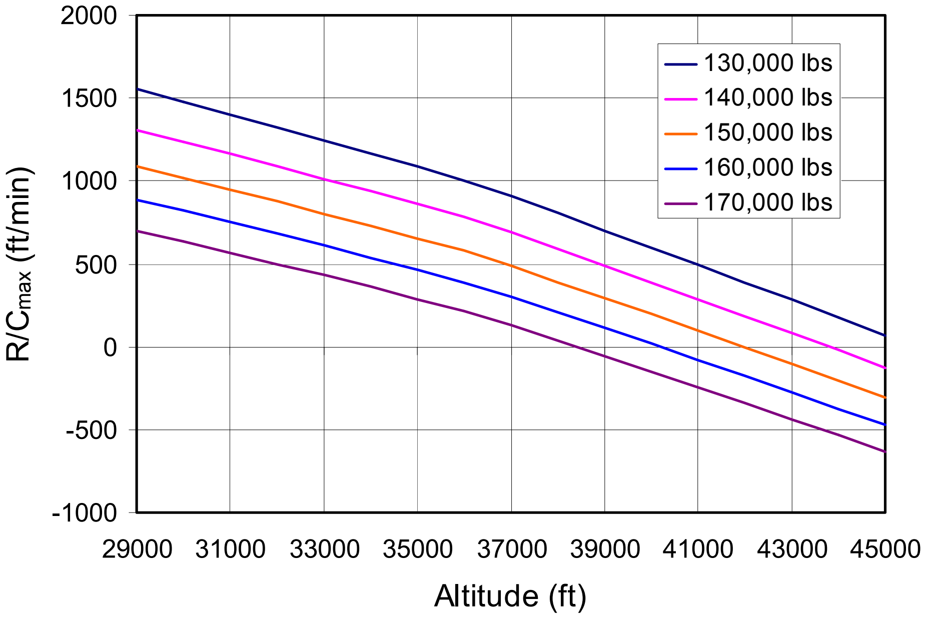

Figure 3.3. The sensitivity of maximum rate of climb to altitude for a range of Mach numbers and weights for a mid-sized jet transport.

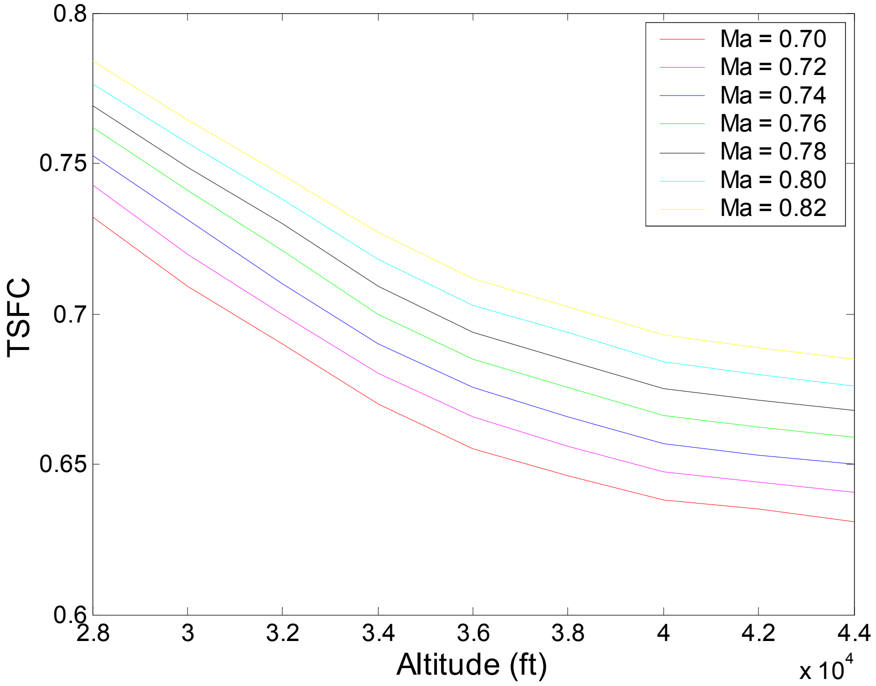

Figure 3.4. Variation of TSFC with altitude for lines of different Mach number.

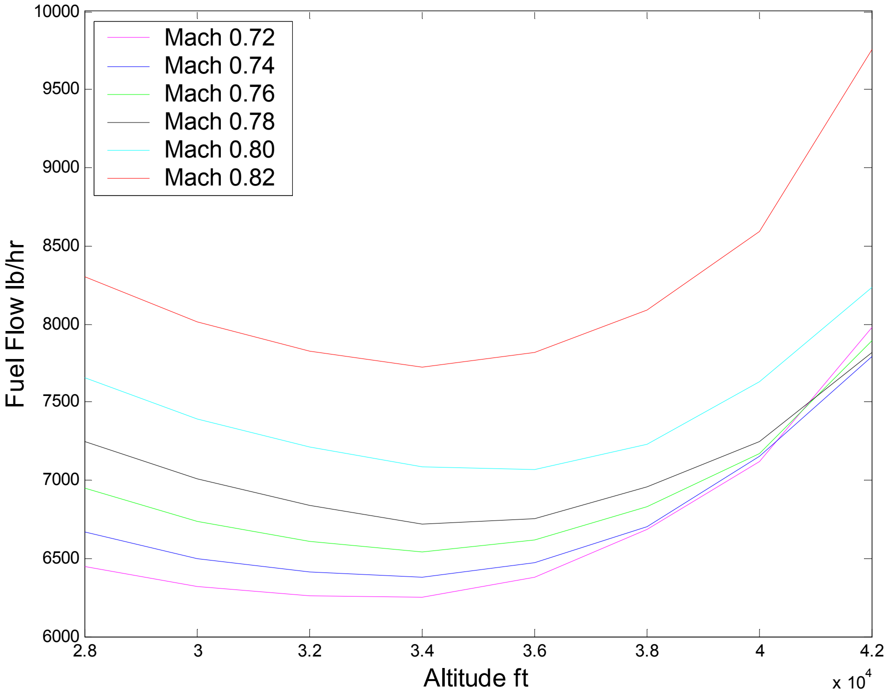

Figure 3.5. Altitude vs. Fuel Flow for W = 145,000 lb.

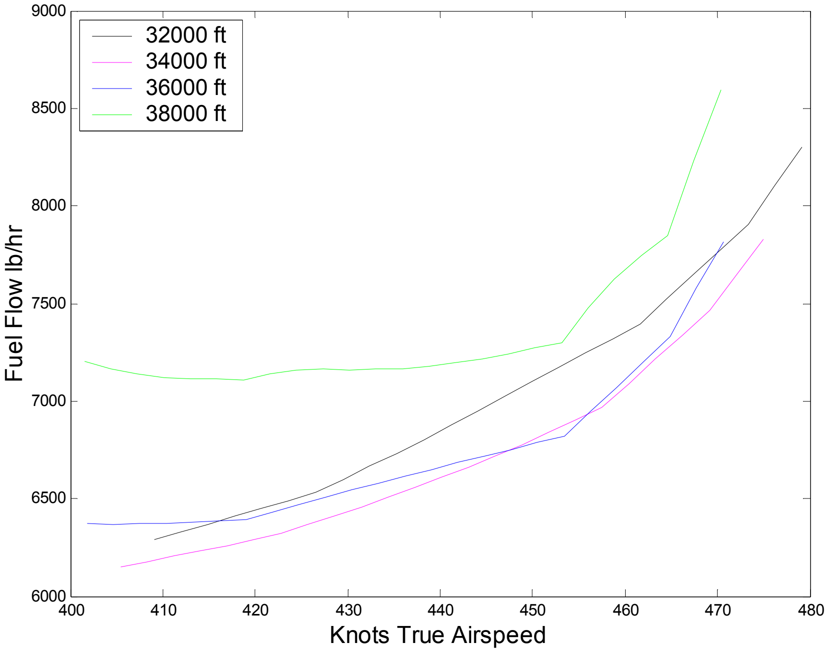

Figure 3.6. Velocity vs. Fuel Flow for W = 145,000 lb.

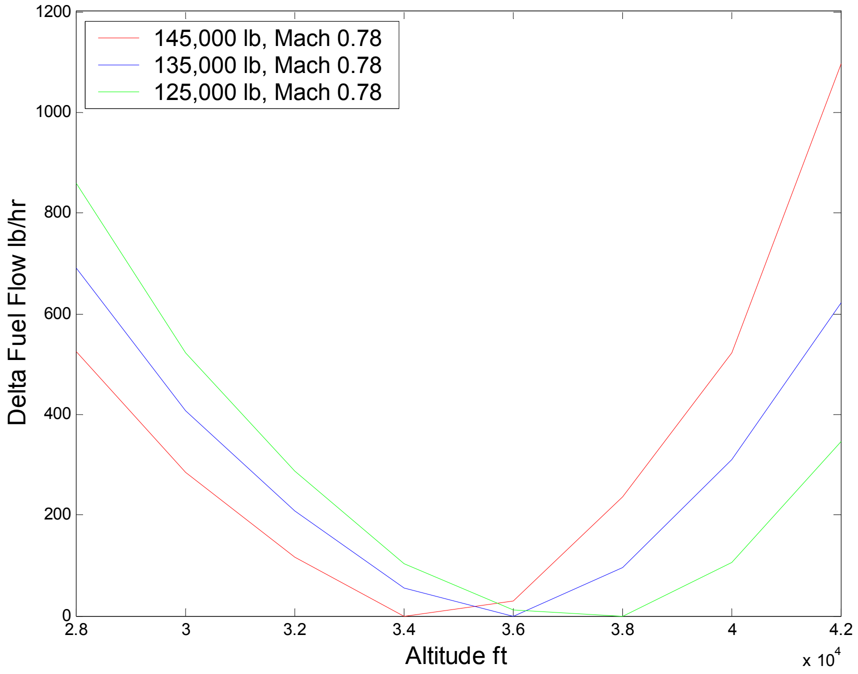

Figure 3.7. Sensitivity of fuel flow to altitude.

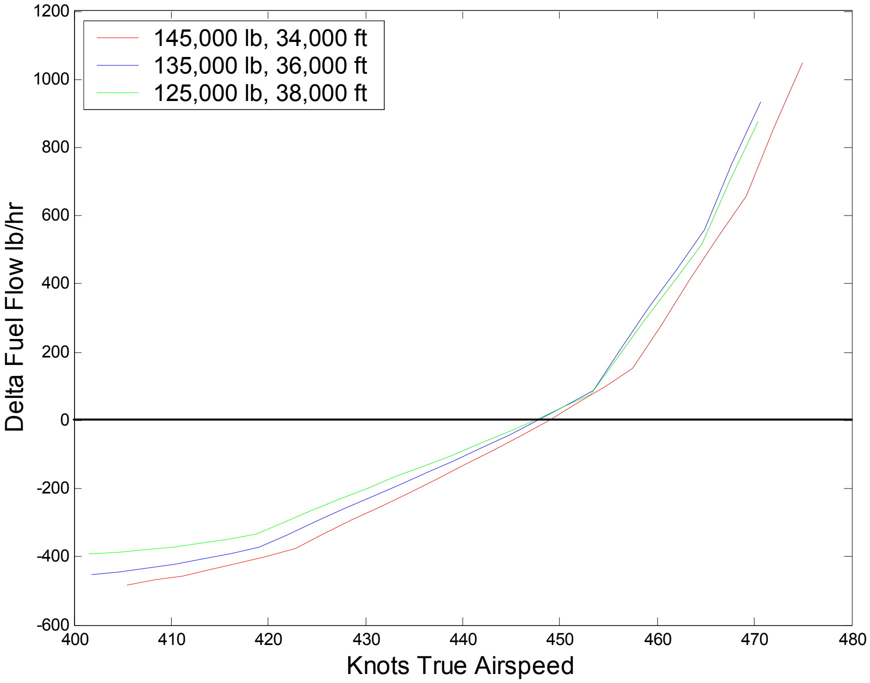

Figure 3.8. Sensitivity of fuel flow to airspeed.

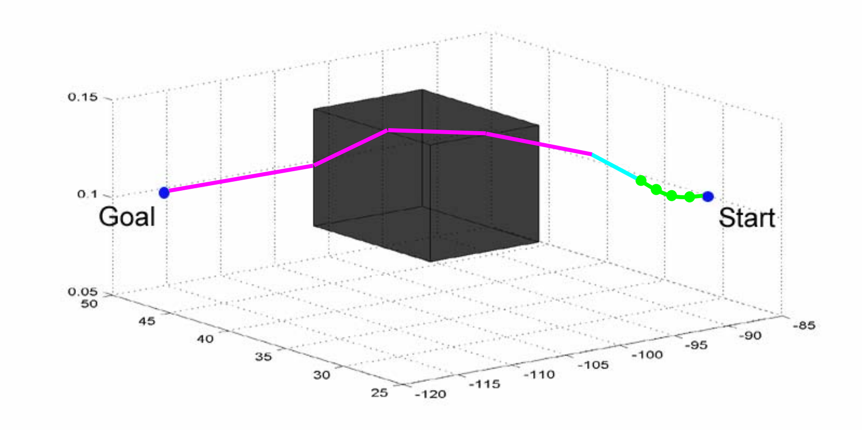

Figure 3.9. Simple example showing the relationship between the detailed trajectory (green), the line of sight (blue), and the cost-to-go (magenta).

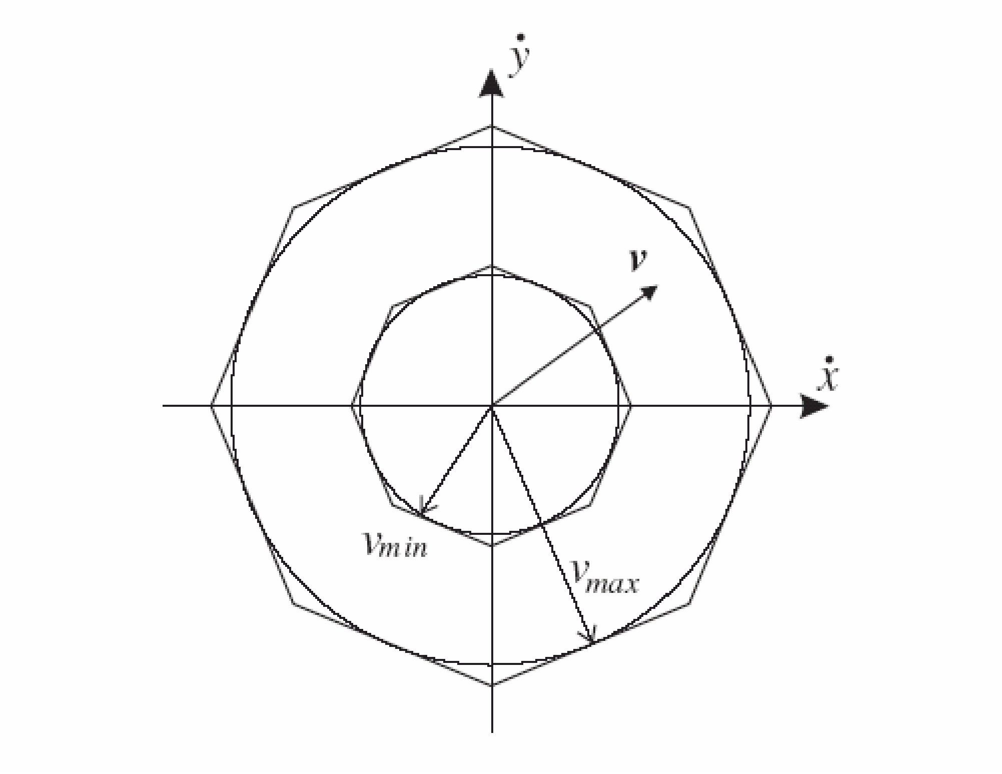

Figure 3.10. Piecewise linear approximation of the 2-norm of the velocity vector [34].

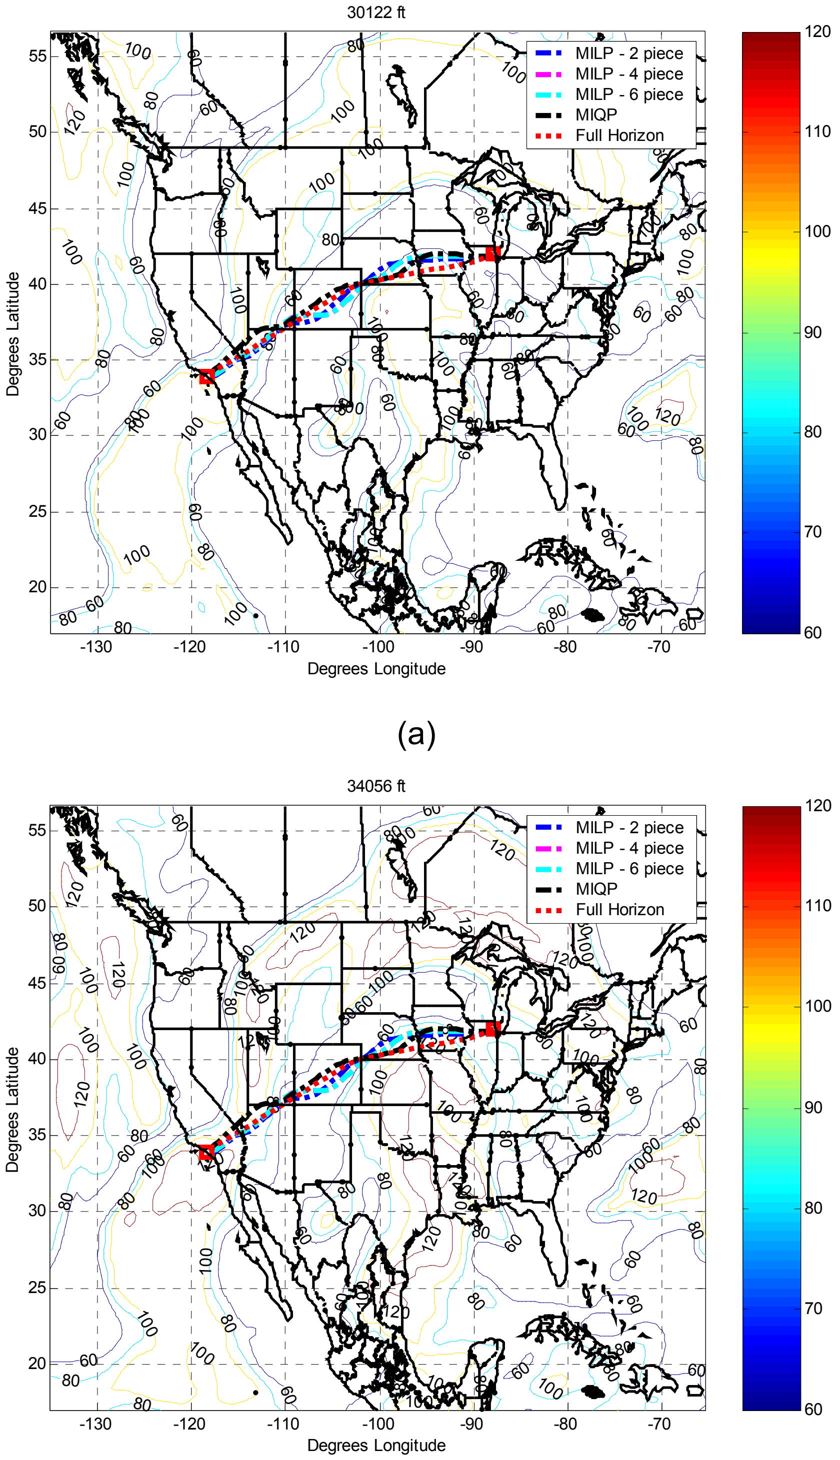

(a)

(b)

Figure 3.11. Fuel optimal trajectories overlaid on contour plots of the RHi field at altitudes of (a) 30122 ft, (b) 34056 ft. A map showing the boundaries of North America is in the background.

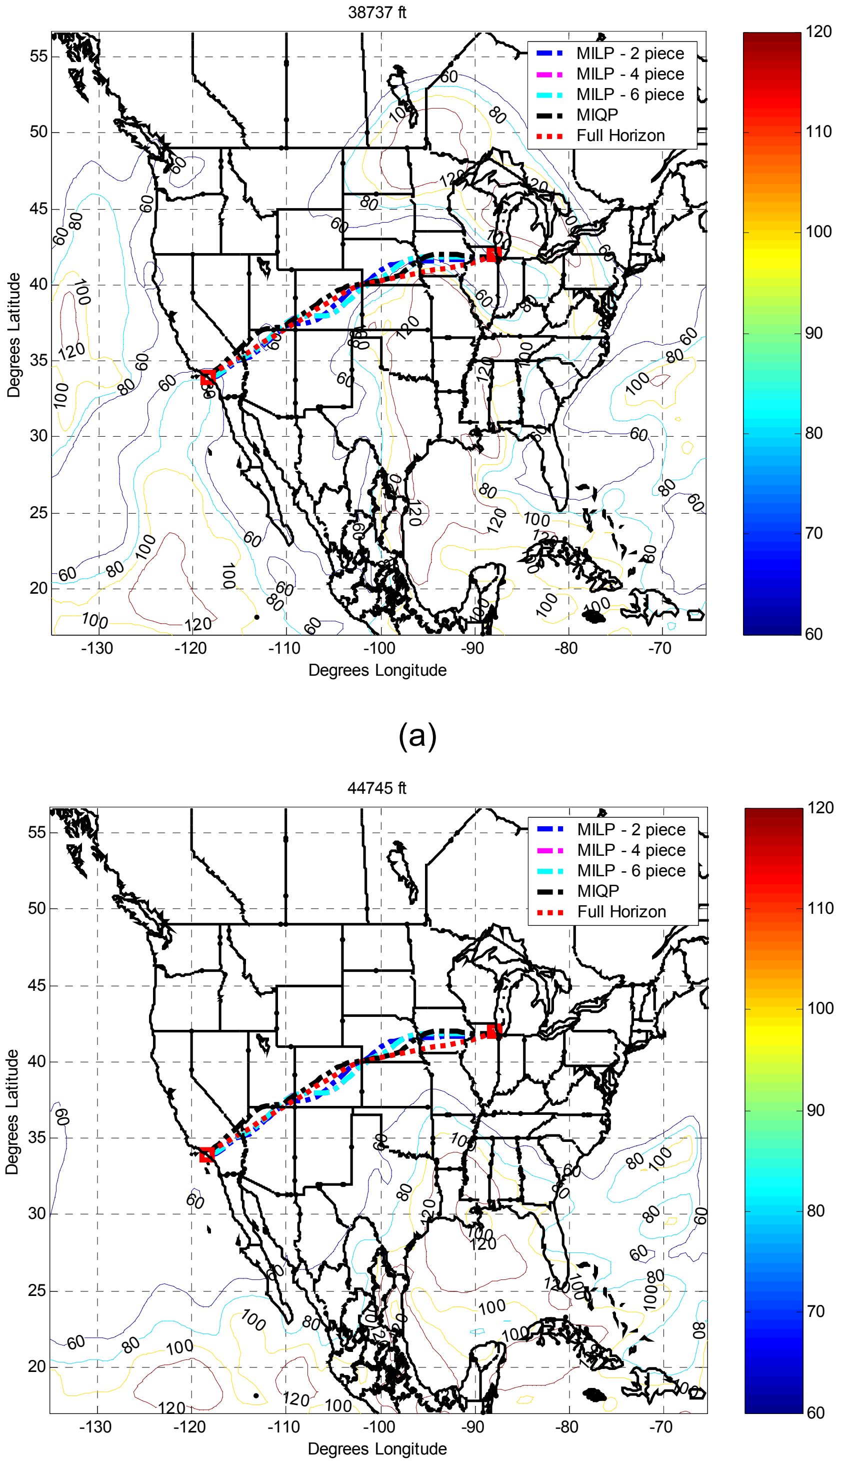

(a)

(b)

Figure 3.12. Fuel optimal trajectories overlaid on contour plots of the RHi field at altitudes of (a) 38737 ft, and (b) 44745 ft. A map showing the boundaries of North America is in the background.

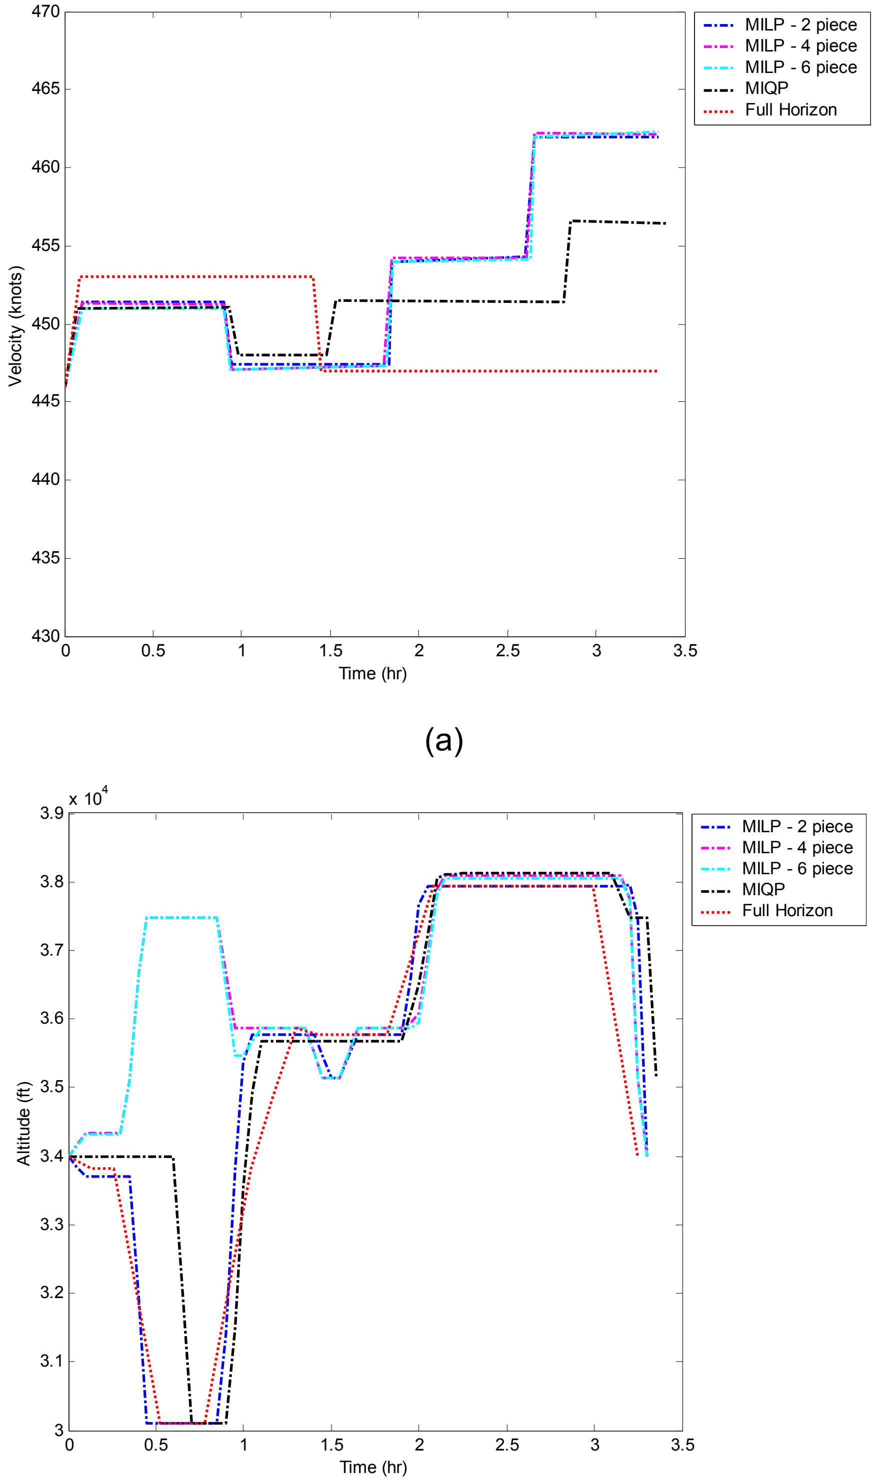

(a)

(b)

Figure 3.13. Aircraft and computational performance: (a) velocity time history, (b) altitude time history.

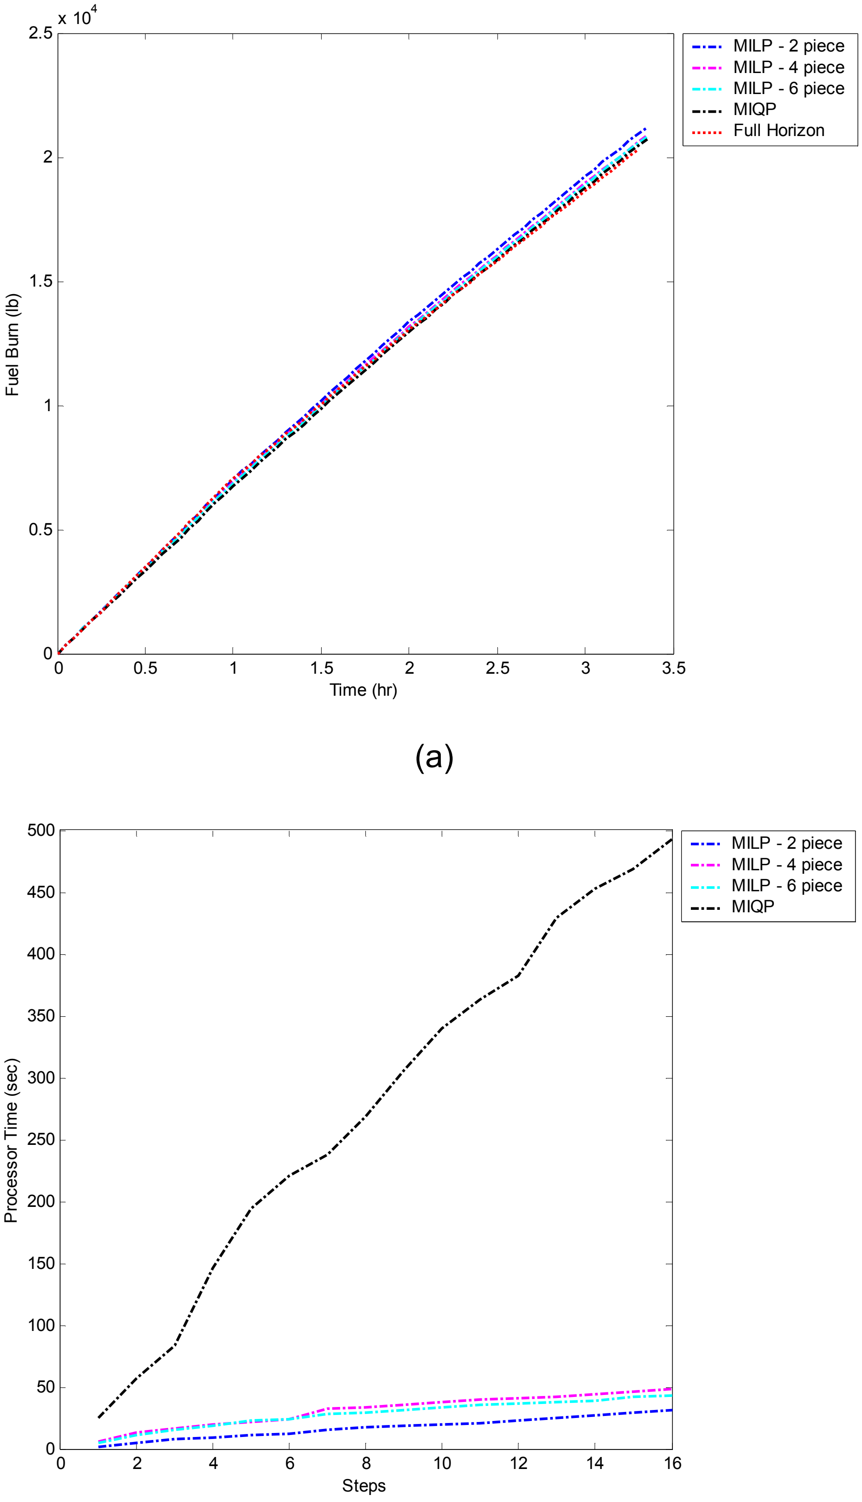

(b)

Figure 3.14. Aircraft and computational performance: (a) fuel burn time history, (b) receding horizon computation time.

Chapter 4

Receding Horizon Mixed-Integer Programming with Soft Avoidance

4.1 Introduction

This chapter presents a new path planning algorithm that is based on the formulation of Chapter 3, but is modified to include soft obstacle avoidance. The phrase "soft obstacle avoidance" means that instead of strictly forcing feasible trajectories outside of obstacles, feasible trajectories are allowed within obstacles, but with an accrued cost penalty. This type of avoidance constraint is useful when obstacle penetration is not a safety of flight issue, such as the case with persistent contrail formation. In the case of persistent contrail mitigation, soft avoidance constraints are attractive because strict avoidance is sometimes infeasible and cost prohibitive from an operational standpoint. For example, consider a large ice super saturated region over the middle of the United States. Hard avoidance would force the flight path completely outside of the region, greatly increasing fuel cost and possibly delaying arrival time. Soft avoidance allows the user to tradeoff fuel and arrival time costs with persistent contrail formation costs, which is much more appealing from an operational standpoint. Consider another example where the airspace above an airport is dominated by an ice super saturated region. Hard avoidance would prevent flights in and out of this airport, which is realistically operationally infeasible.

The cost-to-go function described in Chapter 3, and in the literature [40, 41], forms a path of least cost through the vertices of the obstacles in the environment. This approach is not optimal for a soft avoidance scenario because it forces the trajectory to the boundaries of the obstacle. This chapter presents a novel cost-to-go approximation that allows the trajectory to pass through an obstacle without using the vertices of the obstacle as nodes in the cost grid. Additionally, a new receding horizon algorithm which synthesizes the new cost-to-go and the detailed trajectory is explained and simulation results are presented. The scalability of the new algorithm is simulated and the findings are discussed.

This chapter is organized as follows. First, new avoidance constraints are presented which allow a trajectory to penetrate an obstacle, but with an assigned cost. Next, a novel cost-to-go formulation and receding horizon control approach are explained, and simulation results of the new formulation are discussed. Lastly, the results of the scalability analysis of the algorithm are investigated.

4.2 Obstacle Avoidance and Penalty

The obstacles of this chapter are formulated to be soft, meaning that they can be penetrated, but at a cost. This section describes the formulation of the constraints for soft obstacle avoidance, and highlights the differences between this formulation and that of Chapter 3. Additionally, this section discusses obstacle dynamics, and how the receding horizon controller accounts for the dynamics. Lastly, it should be noted that the dynamical model and fuel cost used in the optimizations of this chapter are identical to the presentation in Chapter 3.

4.2.1 Soft Avoidance Constraints

Persistent contrail mitigation is accomplished within the MILP framework by defining areas conducive to contrail formation as cuboids, and then penalizing flight through these regions, which is solvable by mixed-integer linear programming. Equations 4.1 and 4.2 show the constraints used to penalize persistent contrail formation.

$$\begin{aligned} x[k] &\leq B_{x,low} + R \cdot b_{o,1}[k] \\ x[k] &\geq B_{x,high} - R \cdot b_{o,2}[k] \\ y[k] &\leq B_{y,low} + R \cdot b_{o,3}[k] \\ y[k] &\geq B_{y,high} - R \cdot b_{o,4}[k] \\ z[k] &\leq B_{z,low} + R \cdot b_{o,5}[k] \\ z[k] &\geq B_{z,high} - R \cdot b_{o,6}[k] \end{aligned} \tag{4.1}$$

$$\begin{aligned} &\sum_{i=1}^{6} b_{o,i}[k] \leq 5 + g[k] \\ &g[k] = \{0,1\} \quad for \quad k = 1,\ldots,N_p \end{aligned} \tag{4.2}$$

The variable *B* defines the 6 planes that compose the cuboid, which represents an area of RHi > 100%. The variable $b_{o,i}$ is a binary variable that is either 1 or 0 depending on whether constraint *i* is active, and *g* is a binary variable that takes the value 1 when all 6 constraints are active. This variable is weighted in the cost function and assigns a penalty to contrail formation. The difference between these avoidance constraints and the constraints of Chapter 3.3.3.2 is the binary variable *g[k]* in Eq. 4.2.

4.2.2 Obstacle Dynamics

Obstacle movement and evolution are predicted with a simple model based on the size of the obstacle and its observed movement in the past time step. The general procedure can be described as finding the centroid of each avoidance area, and projecting the centroid of each area forward in time based on how the centroid moved in the last two time steps. For the cuboid representation used in this dissertation, the centroid is found with Eq. 4.3

$$(x_c, y_c, z_c) = \left(\frac{x_{high} - x_{low}}{2}, \frac{y_{high} - y_{low}}{2}, \frac{z_{high} - z_{low}}{2} \right) \tag{4.3}$$

where (x_c, y_c, z_c) is the location of the centroid, and the high and low subscripts refer to the maximum and minimum values of the cuboid in the x, y, and z directions. The velocity of the centroid was found with Eq. 4.4

$$(v_{x,c}, v_{y,c}, v_{z,c}) = \left(\frac{x_{c,i} - x_{c,i-2}}{2dt}, \frac{y_{c,i} - y_{c,i-2}}{2dt}, \frac{z_{c,i} - z_{c,i-2}}{2dt} \right) \tag{4.4}$$

where ($v_{x,c}$, $v_{y,c}$, $v_{z,c}$) is the velocity of the centroid, and *dt* is the size of the time step. The future position of the thunderstorm is simply calculated with Eq. 4.5

$$(x_{c,i}, y_{c,i}, z_{c,i}) = (x_{c,0} + i \cdot v_{x,c} dt, y_{c,0} + i \cdot v_{y,c} dt, z_{c,0} + i \cdot v_{z,c} dt) \tag{4.5}$$

where ($x_{c,i}$, $y_{c,i}$, $z_{c,i}$) is the position of the centroid at time step *i*, and ($x_{c,0}$, $y_{c,0}$, $z_{c,0}$) is the present position centroid.

4.3 Cost-to-Go Formulation

In previous studies involving receding horizon MILP, the cost-to-go was a function of the vertices of the obstacles in the environment [40-44]. This formulation works well when the constraints of the optimization promote hard obstacle avoidance because the optimal trajectory tends to follow along the vertices of the obstacles. However, in the case of soft obstacle avoidance, the optimal trajectory does not generally follow a path through the obstacles in the environment. Therefore, a cost-to-go approximation capable of generating an optimal path through an obstacle was developed. This section presents a novel receding horizon approach and cost-to-go capable of generating nearly optimal trajectories through obstacles with soft avoidance constraints.

4.3.1 Algorithm Overview

The development of this receding horizon control algorithm is based on lessons learned in previous work [25], and is intended to create nearly optimal paths though obstacles described by soft avoidance constraints. The algorithm is formulated through six computational phases, enumerated below and explicitly defined in the following text. The steps to the algorithm are as follows:

1. Find nominal trajectory disregarding persistent contrail formation
2. Create cost grid relative to the aircraft initial position
3. Populate cost grid and create visibility graph
4. Find path of least cost using a modified Dijkstra's algorithm
5. Find optimal trajectory mitigating persistent contrail formation
6. Update initial position and repeat until the destination is reached

The receding horizon algorithm and the cost-to-go associated with the persistent contrail mitigation formulation are described in the following steps:

1. Find nominal trajectory disregarding persistent contrail formation

Existing work has found that trajectories generated with the framework of Chapter 3 tend to be overly conservative when the trajectory is generated in an environment with a limited number of obstacles [25]. The receding horizon algorithm developed in this chapter is intended to improve the optimality by penalizing flight away from a nominal obstacle-free optimal trajectory. The obstacle-free optimal trajectory is solved using identical dynamical constraints, aircraft performance limitations, and fuel burn cost as the formulation of Chapter 3. However, unlike the formulation of Chapter 3, this trajectory is generated with a planning horizon that stretches until the destination, negating the need for a cost-to-go. Normally, a planning horizon of this length would be too much of a computational burden; however, the computation time for this trajectory is significantly reduced due to the lack of obstacles, and corresponding binary variables in the optimization.

2. Create cost grid relative to the aircraft initial position

In order to allow flight through an obstacle, the cost-to-go must include nodes that lie within the obstacles, which is not possible using the framework of Chapter 3. Therefore, a cost grid was created to represent the far-field of the receding horizon controller. The grid consists of sets of nodes that extend radially out from the aircraft, and are stacked vertically to assemble a three-dimensional grid. Figure 4.1 shows a top-down view of the cost grid. It is important to keep in mind that the cost grid is 3-dimensional, and that Fig. 4.1 only represents one layer of the grid. The nodes (black dots in Fig. 4.1) lie outside of the space reachable by the aircraft in the planning horizon and were defined by the following equations:

$$r_0 = V_{\max} \cdot N_p \cdot \Delta t \tag{4.6}$$

$$x_{i,j,k} = (r_0 + \delta_i)\sin\theta_j \tag{4.7}$$

$$y_{i,j,k} = (r_0 + \delta_i)\cos\theta_j \tag{4.8}$$

$$z_{i,j,k} = z_0 + k \cdot \xi \tag{4.9}$$

$$\delta_{i+1} = \delta_i + i \cdot \varepsilon \tag{4.10}$$

$$\theta j = \{\theta_1, \ldots, \theta_n\} \text{ for } j = 1 \ldots n \tag{4.11}$$

where r_0 is the reachable distance in the planning horizon, V_{max} is maximum velocity, N_p is the number of steps in the planning horizon, and Δt is the size of the time step. In Eqs. 4.7-4.9, $x_{i,j,k}$, $y_{i,j,k}$, and $z_{i,j,k}$ are the position coordinates for the nodes, where δ_i is the radial increment of the node away from the airplane, θ_j is the angular increment of the node, z_0 is the altitude of the bottom layer in the grid, and ξ is a user determined constant that spaces the vertical layers of the grid. In Equation 4.10, ε is a user determined constant that spaces the radial increments of the grid. The size of the grid is set by the user prior to the execution of the algorithm.

3. Populate cost grid and create visibility graph

The cost-to-go is determined by finding a path of least cost through the cost grid created in step 2. Each node in the cost grid is assigned a value based on the distance from the nominal trajectory and whether or not the node lies in an area of ice supersaturation. The distance from the nominal trajectory is calculated with a linear program, which is given as Eq. 4.12.

$$\begin{aligned} &\min \quad c^T D_{i,j,k} \\ &\text{subject to}: D_{\mathrm{i,j,k}} \geq \sqrt{\left(x_{i,j,k} - x_{nom,l}\right)^2 + \left(y_{i,j,k} - y_{nom,l}\right)^2 + \left(z_{i,j,k} - z_{nom,l}\right)^2} \end{aligned} \tag{4.12}$$

where $D_{i,j,k,l}$ is a vector of the distances from each node in the cost grid to the nominal trajectory, c is a n x 1 vector of 1's, where n is the number of nodes in the grid, ($x_{i,j,k}$, $y_{i,j,k}$, $z_{i,j,k}$) are the position coordinates for the nodes, and ($x_{nom,l}$, $y_{nom,l}$, $z_{nom,l}$) are the position coordinates of the l^{th} waypoint in the nominal trajectory. The total cost of each node is defined by Eq. 4.13

$$\Phi_{i,j,k} = \gamma \cdot D_{i,j,k} + \upsilon \cdot P_{i,j,k} \tag{4.13}$$

where γ is a user defined weighting on the distance from the nominal trajectory, υ is a user defined weighting on persistent contrail formation, and $P_{i,j,k}$ is a vector of 1's and 0's where a 1 indicates that the node lies within an area of ice supersaturation.

4. Find path of least cost using a modified Dijkstra's algorithm

As mentioned previously, the cost-to-go is the path of least cost through the cost grid. The path of least cost is found using a modified version of Dijkstra's algorithm, where Dijkstra's algorithm is a method used to find the shortest path through a graph of nodes [76]. This step modifies Dijkstra's algorithm to account for the cost of each node found in step 3 in addition to the distance between the nodes, which means that the algorithm is finding the path of least cost, not the path of least distance. For example, the cost to travel between nodes (1,1,1) and (2,2,2) is given by Eq. 4.14

$$\Psi_{(1,1,1),(2,2,2)} = \Phi_{1,1,1} + \Phi_{2,2,2} + \rho \cdot \sqrt{\left(x_{1,1,1} - x_{2,2,2}\right)^2 + \left(y_{1,1,1} - y_{2,2,2}\right)^2 + \left(z_{1,1,1} - z_{2,2,2}\right)^2} \quad (4.14)$$

where ρ is a user defined weighting on the distance between nodes. The path of least cost is shown in Fig. 4.2 as the magenta line passing through the cost grid. The aircraft in Fig. 4.2 is traveling from right to left, where the black crosses show the executed trajectory, the green crosses show the planning horizon, the blue line connects the planning horizon to the cost grid, and the area of ice supersaturation is depicted by the gray box. Note that this figure is shown in two dimensions for ease of understanding, and that the simulation is actually in three dimensions.

5. Find optimal trajectory mitigating persistent contrail formation

The trajectory is optimized in a receding horizon fashion where the total cost of the trajectory is the sum of the cost in the planning horizon, the cost-to-go, and the line of sight, which connects the end of the planning horizon to the cost grid. The total cost of the trajectory is shown in Eq. 4.15

$$J_{total} = \sum_{i=1}^{N_p} \left(f_{a,i} + f_{v,i} + A_i + b_{f,i} f_{term,i} + b_{climb,i} f_{descend} + b_{descend,i} f_{descend} \right) + \Psi_{cost-to-go} + \eta \cdot D_{LOS} \quad (4.15)$$

where $\Psi_{cost\text{-}to\text{-}go}$ is the cost-to-go, D_{LOS} is the distance between the end of the planning horizon and the start of the cost-to-go, η is a weighting on the line of sight, and the rest of the equation is the cost in the planning horizon. It should be reminded that the following weightings have an effect on the behavior of the trajectory: the nominal trajectory and persistent contrail weighting (Eq. 4.13), the node distance weighting (Eq. 4.14), and the line of sight weighting (Eq. 4.15). In the cases of the node distance and line of sight weightings, the weighting factors were scaled to approximate an average fuel burn over the given distance based on nominal aircraft performance parameters. For example, the line of sight weighting was approximated with Eq. 4.16

$$\rho \approx \frac{W_{f,cruise}}{V_{cruise}} \tag{4.16}$$

where $W_{f,cruise}$ is the nominal fuel burn and V_{cruise} is the nominal cruise airspeed. The persistent contrail weighting is treated as a user defined value which determines the likelihood of persistent contrail formation.

6. Update initial position and repeat until the destination is reached

The receding horizon controller executes steps 1-5 of this algorithm, updates the initial position, and then repeats steps 1-5 until the destination is reached. In addition to updating the initial position, the environment is updated to account for any change in the areas of ice supersaturation and to adjust the nominal flight condition of the aircraft, which depends on the aircraft weight at that iteration.

4.3.2 Simulation Results

The algorithm was tested through simulation to assess its viability and to compare its performance to a non-receding horizon MILP trajectory, which is the most optimal solution the method can provide. Figure 4.3 shows the results for an environment containing one large static obstacle. The trajectory was generated with a planning horizon of 16, an execution horizon of 8, and a time step of 2.4 minutes. The penalty for obstacle penetration is large to help illustrate the operation of the algorithm. The obstacle shaded in yellow, the trajectory originates on the right side of the figure, and the goal is located in the bottom left corner. The executed waypoints in the trajectory are identified by black crosses, the planning horizon by green crosses, the line of sight by a blue line, and the cost grid by red circles. The circles in the cost grid are blue when they are located within the boundaries of an obstacle. The receding horizon controller initially recognizes the obstacle in the cost grid, as is observed by the blue circles within the obstacle in Fig. 4.3(a). The cost-to-go finds a path of least cost through the cost grid and around the obstacle, and the receding horizon controller iterates until the goal is reached. In Fig. 4.3(f), the overall shape of the receding horizon trajectory is similar to the non-receding horizon trajectory. Additionally, the receding horizon trajectory performs well compared to the non-receding horizon controller, burning only 0.9% more fuel and arriving 2 minutes later.

Figure 4.4 presents an example with an environment containing three dynamic obstacles. The trajectory was generated with a planning horizon of 16, an execution horizon of 8, and a time step of 2.4 minutes. The penalty for obstacle penetration is large to help illustrate the operation of the algorithm. The obstacles are shaded in red, and the trajectory originates on the right side of the figure and the goal is located in the bottom left corner. The executed waypoints in the trajectory are identified by black crosses, the planning horizon by green crosses, the line of sight by a blue line, and the cost grid by red circles. Compared to the first example, the obstacles in this example are small and do not have the same effect on the cost-to-go. The cost-to-go avoids the obstacles, but it does so without significantly changing the trajectory. As the trajectory progresses, the obstacles become within reach of the planning horizon, and the trajectory begins to change direction, as is seen in Fig. 4.4(c). In Fig. 4.4(d), it is apparent the receding horizon controller predicts the dynamics of the obstacle, allowing the planning horizon to find a feasible solution behind the first obstacle. The receding horizon trajectory continues until the aircraft reaches the destination. Similarly to the previous example, this receding horizon trajectory closely matches the non-receding horizon trajectory, as is seen in Fig. 4.4(f). The receding horizon trajectory burned 0.6% more fuel and arrived 4 minutes later than the non-receding horizon trajectory.

4.4 Scalability of Algorithm

It is important to investigate how the running time of an algorithm scales with the size of the problem. Typically, this analysis is conducted by summing of the running times of each statement executed in the algorithm. Sometimes, an algorithm can be too complex for a simple counting type analysis. In this case, the running time is evaluated through stochastic simulation, where the algorithm is executed with different input sizes and the results are graphed. A scalability analysis of the algorithm developed in this chapter is performed for each step in the algorithm, and then the algorithm is evaluated as a whole.

The effect of nominal trajectory size, cost grid size, and planning horizon size are analyzed for each step. The nominal trajectory size is equal to the number of waypoints in the nominal trajectory. The number of binary variables in the problem depends on the number of waypoints in the nominal trajectory due to the quadratic approximation for velocity, as given in Eqs. 3.11 and 3.12. The cost grid size is the number of nodes in the cost grid. The planning horizon size is equal to the number of waypoints in the planning horizon. The number of binary variables in the problem increases with the planning horizon size because of the quadratic approximation for velocity and the obstacle constraints.

Step 1 of the algorithm is to find a nominal trajectory in absence of obstacles using MILP. It is an analytically ambiguous relation between the operational executions involved in Step 1 and the overall computational time achieved, therefore an experimental analysis is performed. Figure 4.5 shows the sensitivity of Step 1 running time to changes in planning horizon size, nominal trajectory size, and cost grid size. The data points in Fig. 4.5 are averages of approximately 20 trials of the algorithm. It is apparent from looking at Fig. 4.5 that the size of the nominal trajectory has a dominant effect on the running time of Step 1. The effect of doubling the number of waypoints in the nominal trajectory roughly doubles the running time of Step 1. However, it should be noted that this trend should not be expected for a large number of waypoints due to the fact that MILP path planning problems have been found to be NP-Hard.

Step 2 is the creation of the cost grid for the cost-to-go. The construction of the cost grid is relatively straight forward, where the number of statements is equal to the size of the grid. Therefore, we say that Step 2 running time is $O(C_{ij})$, where C_{ij} is the number of nodes in the cost grid and O is the standard “big-O” notation. In mathematics, “big-O” notation describes the limiting behavior of a function when the argument tends towards a particular value or infinity.

Step 3 of the algorithm populates the cost grid and creates a visibility graph for the nodes in the grid. Like Step 1, this step in the algorithm was evaluated via experiment. Figure 4.6 presents the sensitivity of Step 3 running time to changes in planning horizon size, nominal trajectory size, and cost grid size. The data points in Fig. 4.6 are averages of approximately 20 trials of the algorithm. The dominant input on the running time of Step 3 is the size of the cost grid, with the size of the nominal trajectory having a secondary effect. Doubling the size of the cost grid roughly doubles the running time of Step 3. The nominal trajectory size has a small effect on the running time because the distance from each node in the cost grid to the nominal trajectory is computed with a linear program, of which the nominal trajectory size is an input. Linear programming problems have been shown to be solvable in polynomial time.

Step 4 of the algorithm uses a modified version of Dijkstra's algorithm to find the path of least cost through the cost grid. It has been shown that the running time for Dijkstra's algorithm is $O(E^2)$, where E is a list of the vertices in the graph [76]. In the case of Step 4, the running time is therefore $O(C_{ij}^2)$, where C_{ij} is the number of nodes in the cost grid.

Step 5 is the optimization of the detailed trajectory. The running time of this step is evaluated experimentally. Figure 4.7 gives the sensitivity of Step 5 running time to changes in planning horizon size, nominal trajectory size, and cost grid size. The data points in Fig. 4.7 are averages of approximately 20 trials of the algorithm. The dominant effect on the running time is the size of the planning horizon. Doubling the planning horizon roughly doubles the running time as long as the planning horizon is small. For larger planning horizons, doubling the planning horizon size has a much more significant effect of the running time.

An analysis of the complete algorithm is important to show that the optimization can remain tractable with increasing problem size. Mixed-integer linear programming problems have been proven to be NP-Hard in the number of binary variables [46, 47]. The number of binary variables depends primarily on the number of waypoints in the planning horizon and the number of obstacles in the environment. As the number of binary variables in the problem increases, the computation time of the increases exponentially. Figure 4.8 shows the average computation time of the algorithm as the number of binary variables increases. The computation time shows a gradual and consistent increase until approximately 2300 binary variables, at which point the average computation time increases dramatically and becomes unpredictable.

Figure 4.9 shows how the maximum computation time of the algorithm depends on the number of binary variables in the problem. As expected, the maximum computation time increases considerably at approximately 2300 binary variables. Additionally, it is of note that the maximum computation time at approximately 2400 binary variables is 250 seconds, which is large enough to make the path planner intractable for any real time embedded aviation applications. Figure 4.10 presents the dependence of the standard deviation of the computation time to the number of binary variables in the problem. The maximum of the standard deviation occurs at approximately 2500 binary variables and takes a value of 120 seconds.

The algorithm does not have predictable scaling behavior above a certain number of binary variables, which is observed to be approximately 2400. When the problem size reaches this point, the tractability of the algorithm in a real time setting comes into question, due to the significantly increasing average computation time and more importantly, the standard deviation of the computation time. Therefore, the algorithm analysis establishes a bound of 2400 binary variables to help the algorithm maintain tractability for real time applications. This number of binary variables corresponds to approximately 20 obstacles in the environment with a planning horizon of 16 waypoints.

4.5 Summary

This chapter presented a novel receding horizon algorithm and cost-to-go formulation to generate trajectories that obey soft obstacle avoidance. The significance of soft avoidance is that strict avoidance is sometimes infeasible and cost prohibitive from an operational standpoint. This is the case in persistent contrail mitigation, where areas of ice supersaturation can be extremely large and possibly cover the origin and destination of a flight. The steps of the algorithm were discussed and simulations of the algorithm were conducted to investigate its performance, and the receding horizon algorithm performed very well compared to the non-receding horizon formulation. Lastly, the scalability of the algorithm was explored through an empirical investigation.

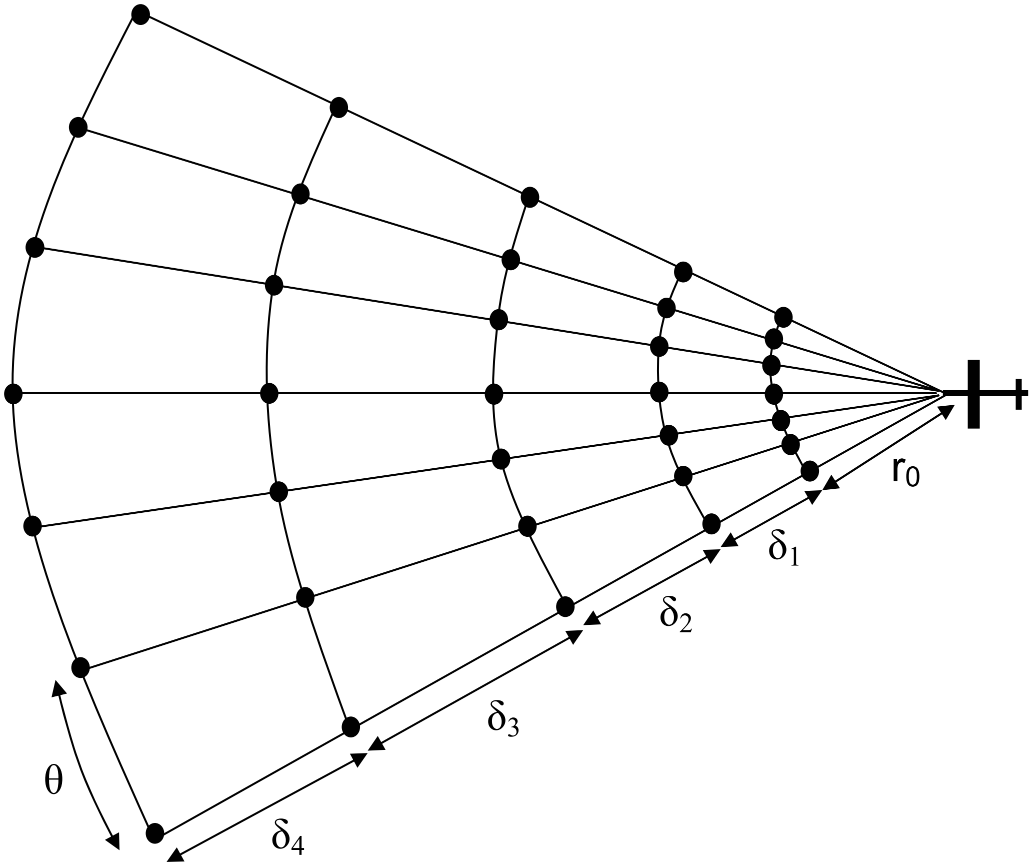

Figure 4.1. Top-down View of Cost Grid. This figure shows one layer of the cost grid.

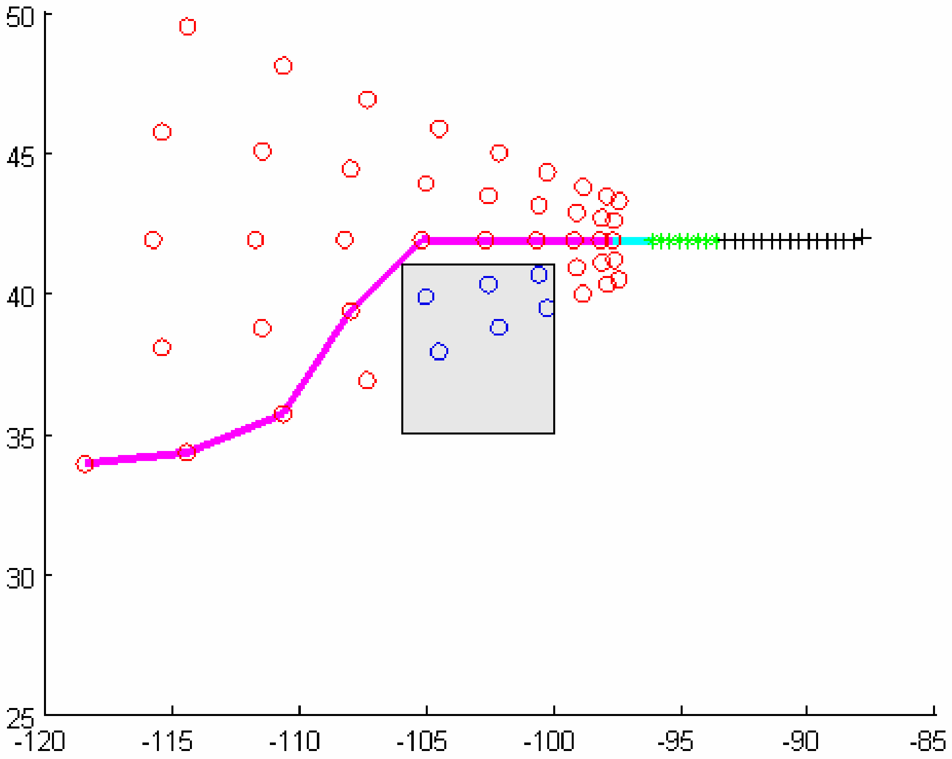

Figure 4.2. Top-down view of trajectory showing the planning horizon (green), line of sight (blue), and the cost-to-go (magenta). The open circles are nodes in the cost grid, and blue circles indicate the node is in an area of ice super saturation (gray box). The black crosses show executed waypoints.

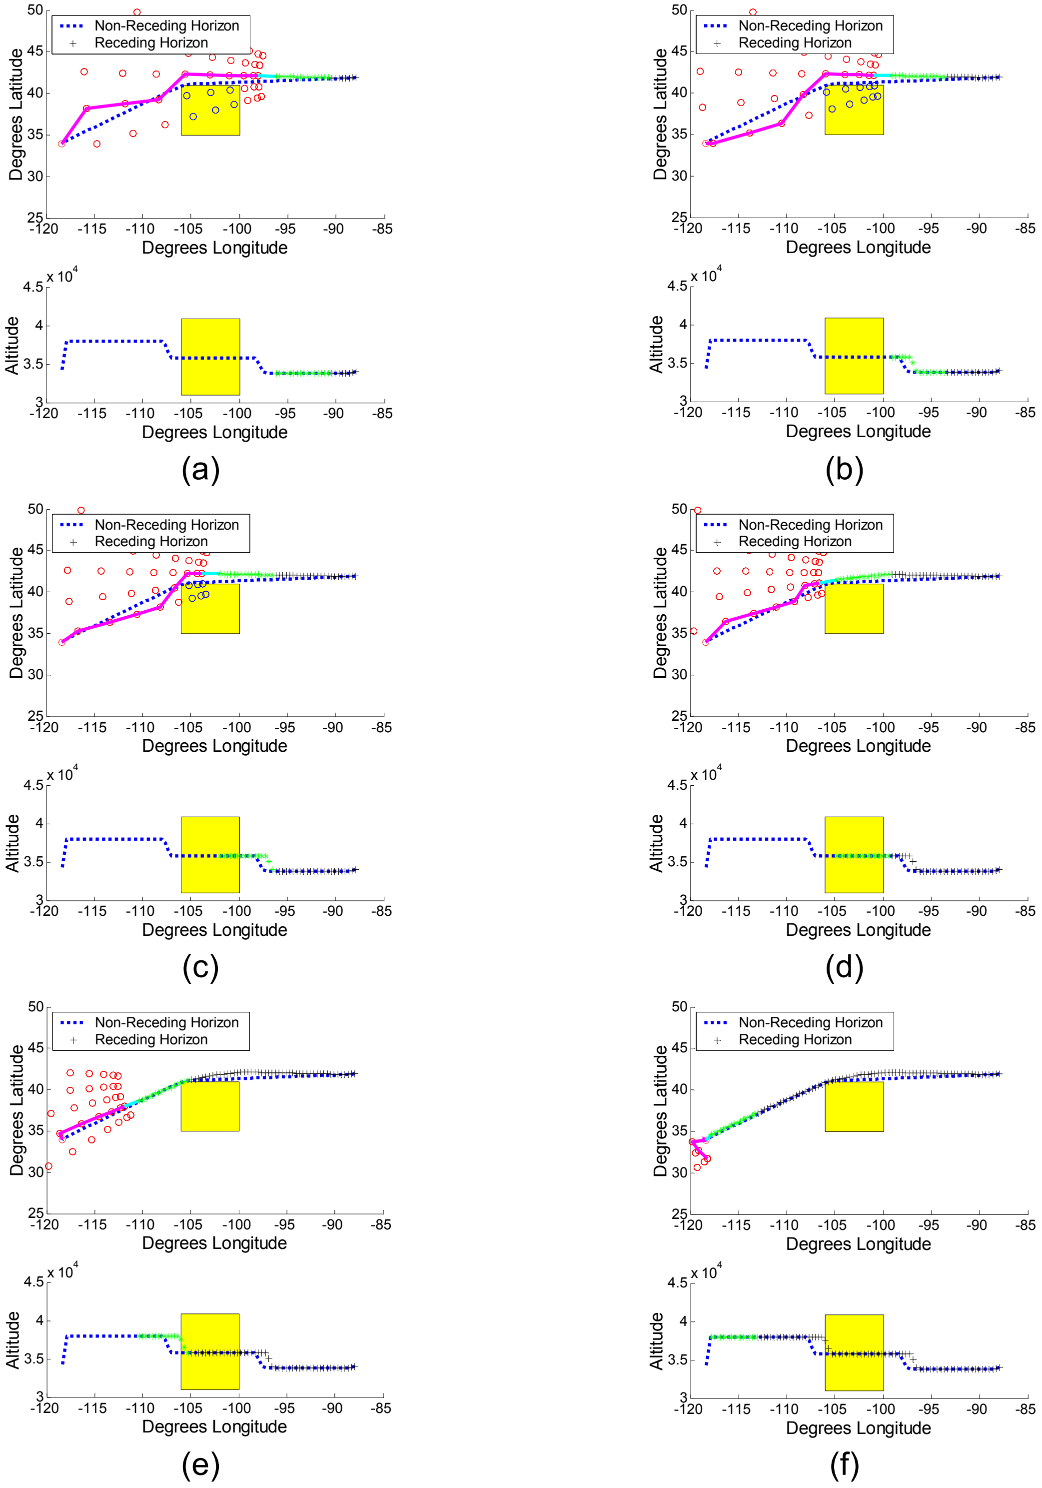

Figure 4.3. Example of receding horizon trajectory around a large sized obstacle. The executed portion of the receding horizon trajectory is shown in black, the planning horizon is in green, and the cost-to-go is magenta. The non-receding horizon trajectory is depicted as a blue dashed line.

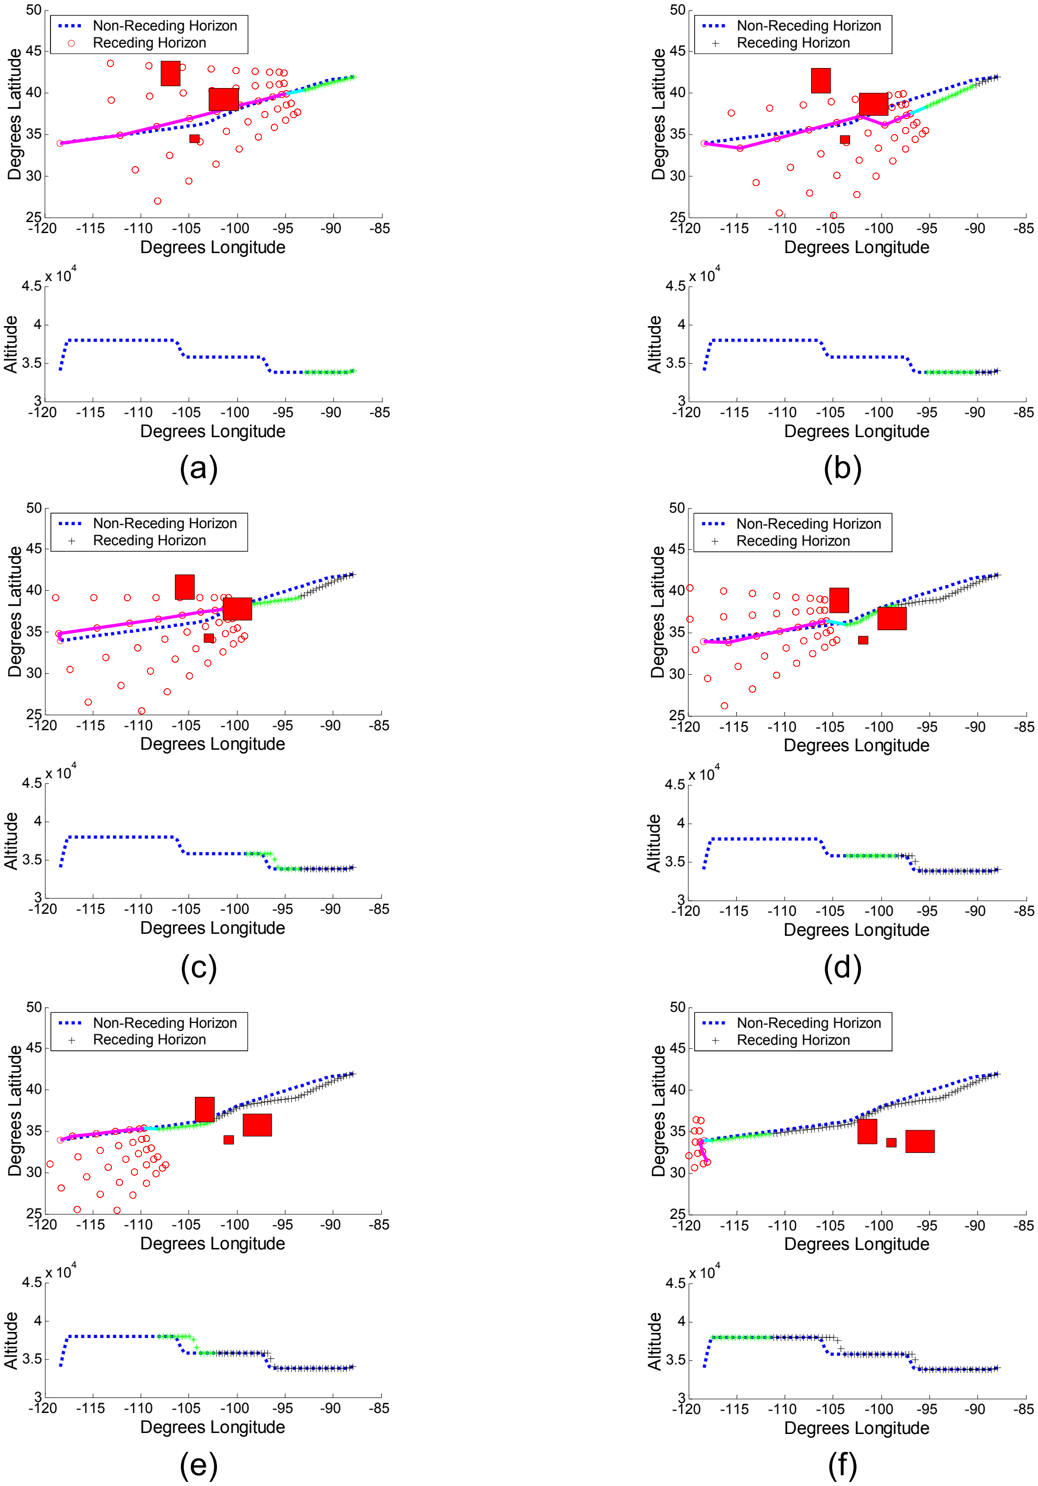

(a) (b) (c) (d) (e) (f)

Figure 4.4. Example of a receding horizon trajectory around three dynamic obstacles. The executed portion of the receding horizon trajectory is shown in black, the planning horizon is in green, and the cost-to-go is magenta. The non-receding horizon trajectory is depicted as a blue dashed line.

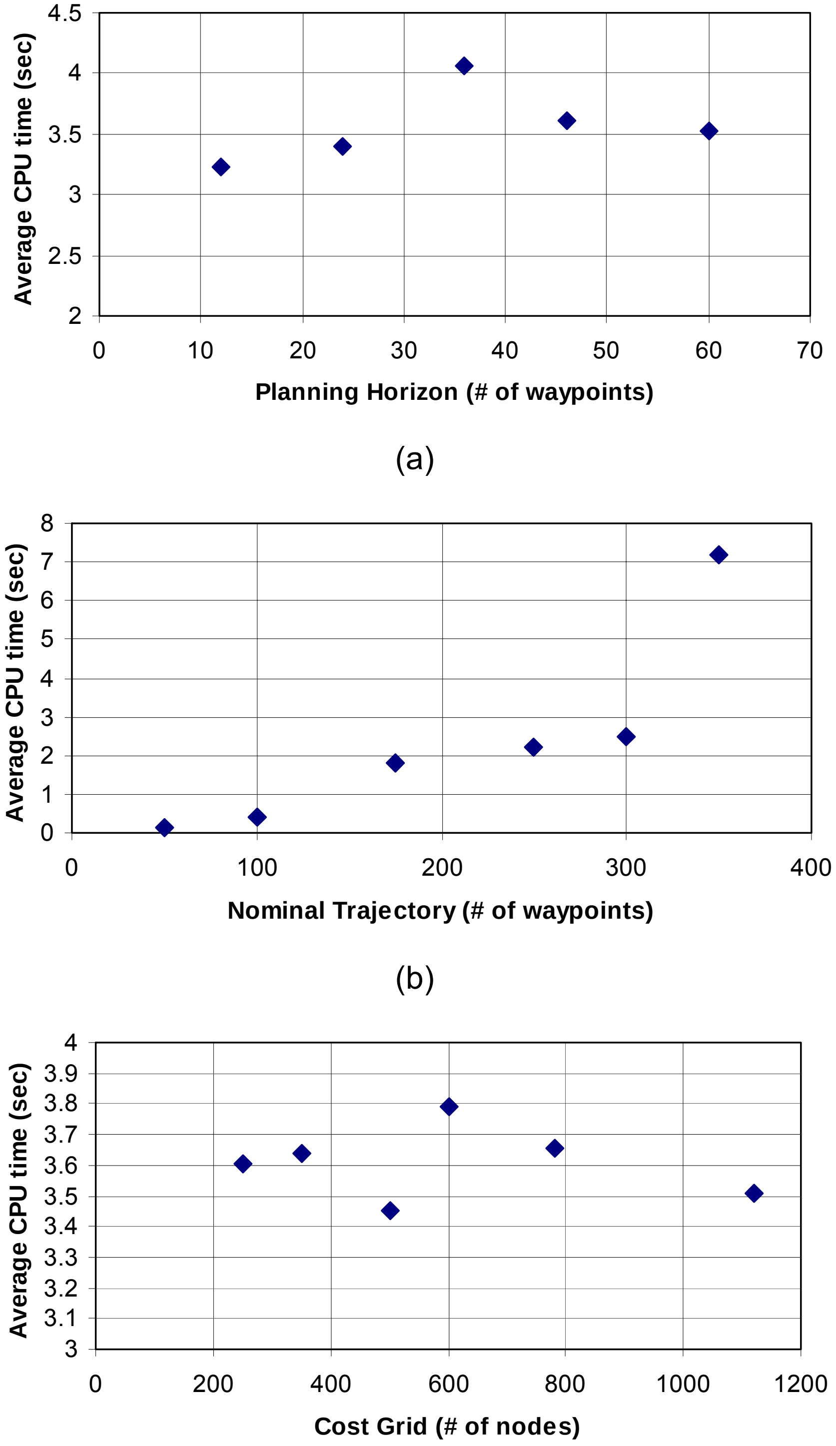

(a)

(b)

(c)

Figure 4.5. Sensitivity of Step 1 running time to (a) planning horizon size, (b) nominal trajectory size, and (c) cost grid size.

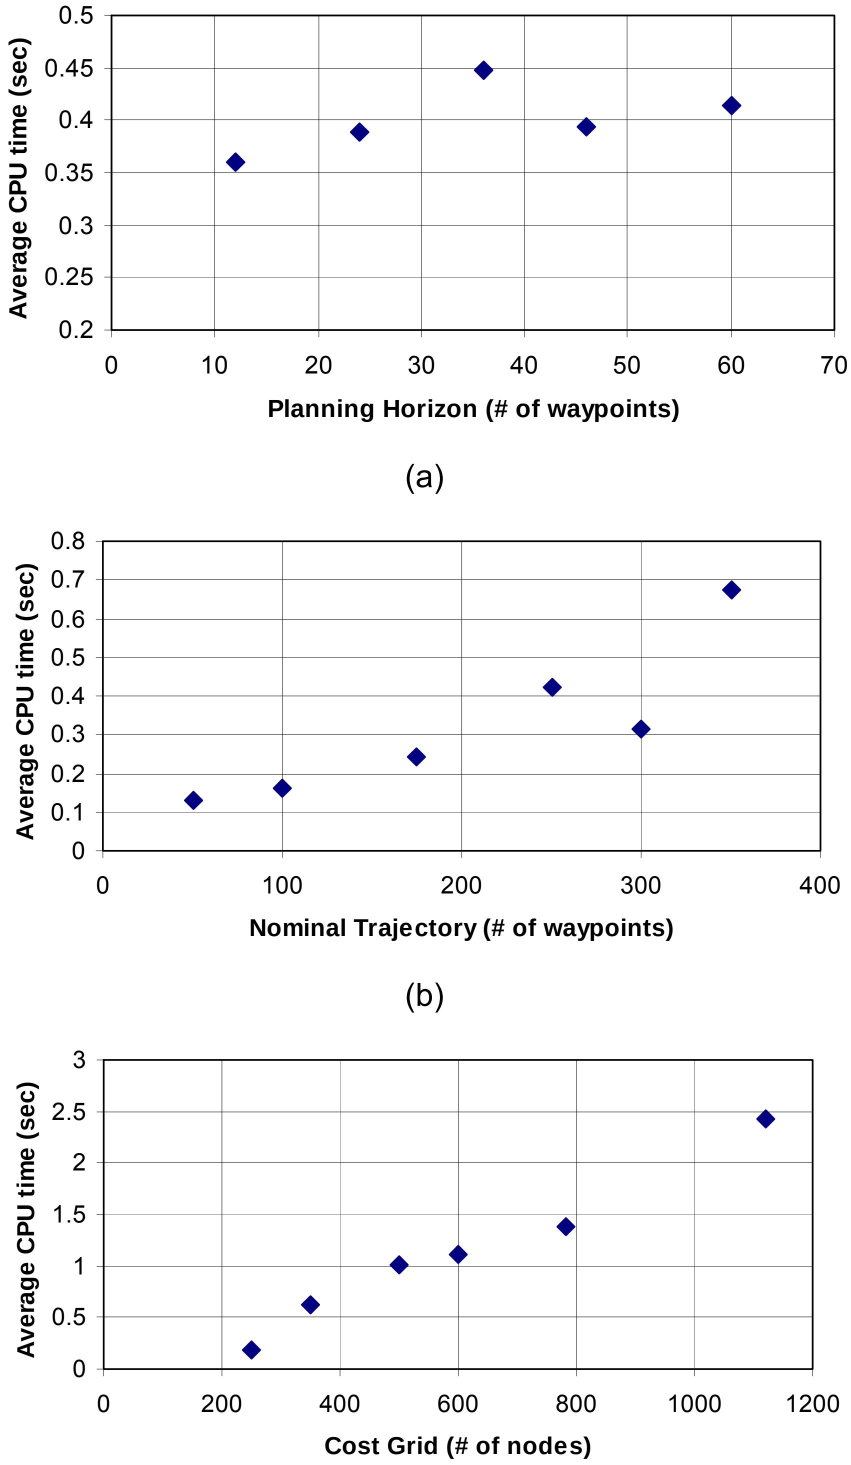

(a)

(b)

(c)

Figure 4.6. Sensitivity of Step 3 running time to (a) planning horizon size, (b) nominal trajectory size, and (c) cost grid size.

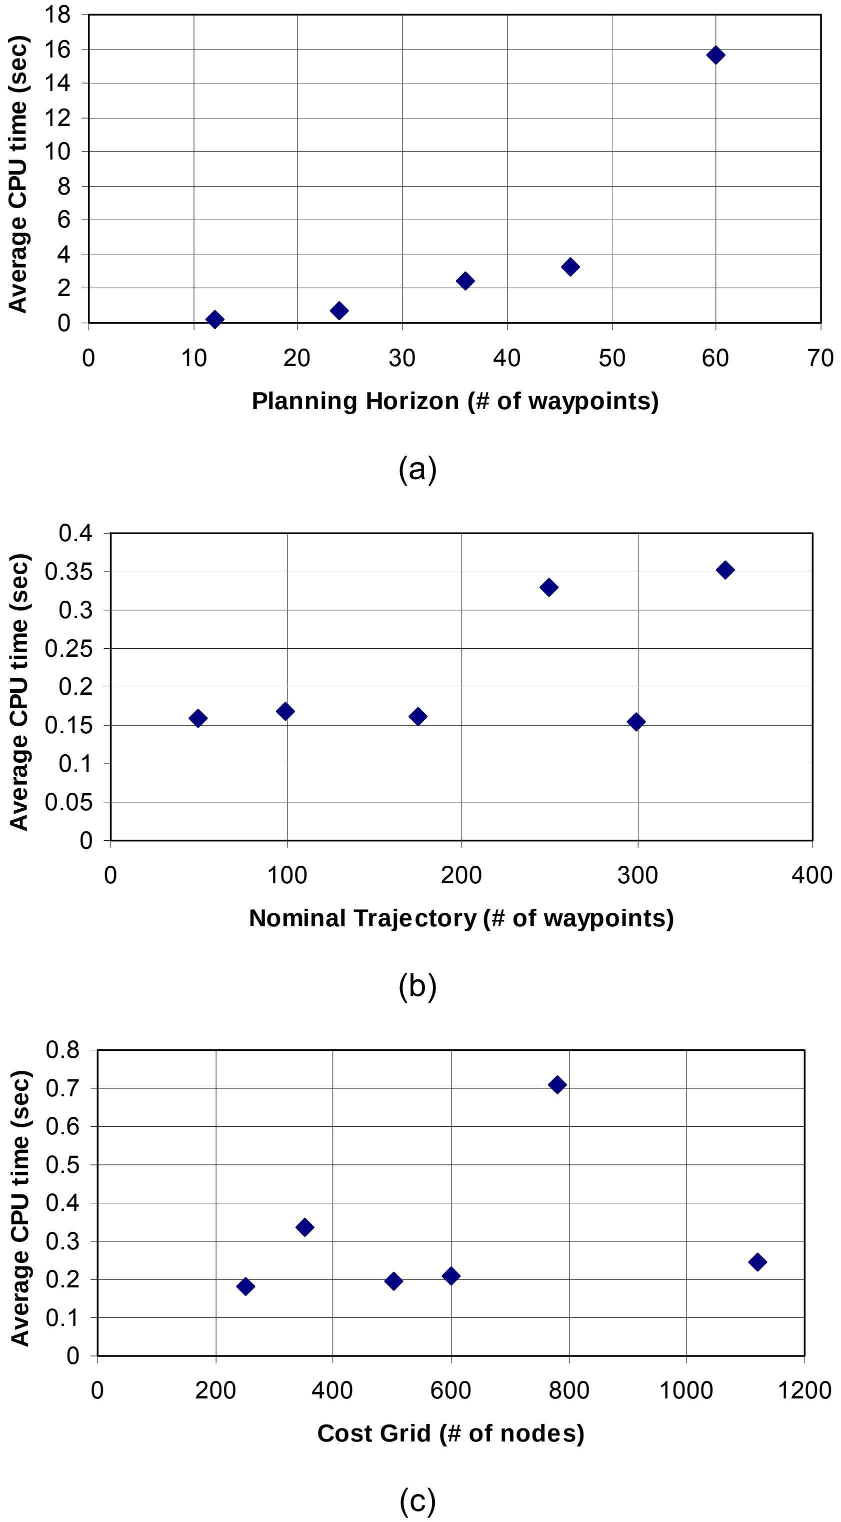

Figure 4.7. Sensitivity of Step 5 running time to (a) planning horizon size, (b) nominal trajectory size, and (c) cost grid size.

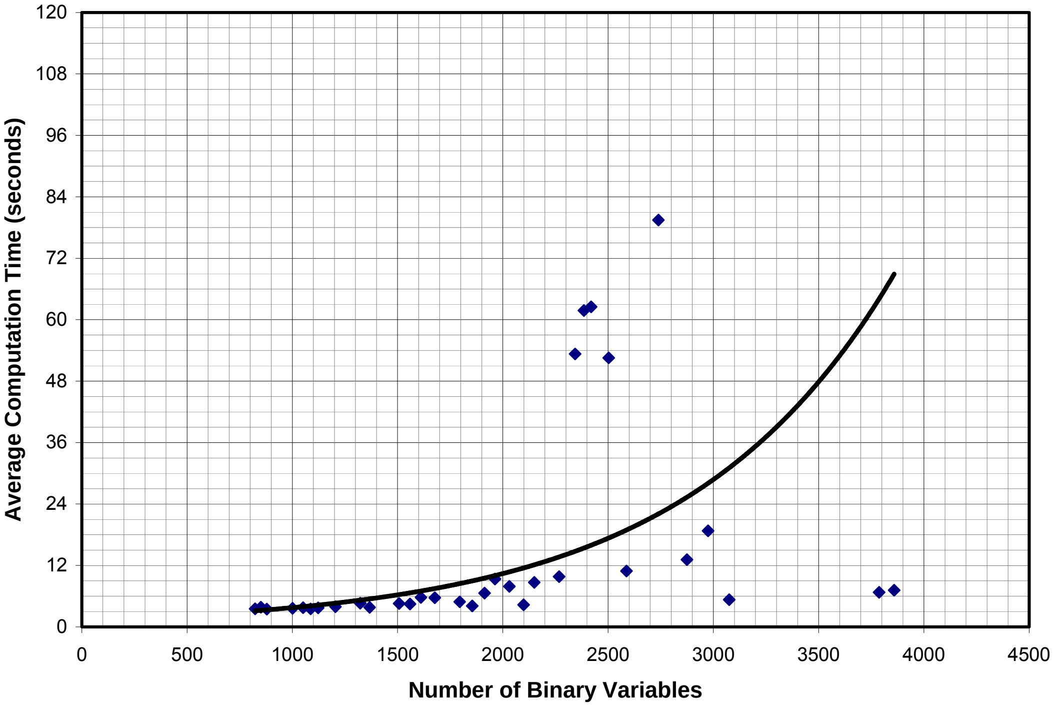

Figure 4.8. Average computation time of algorithm for various problem sizes. The solid black line is an exponential trend line.

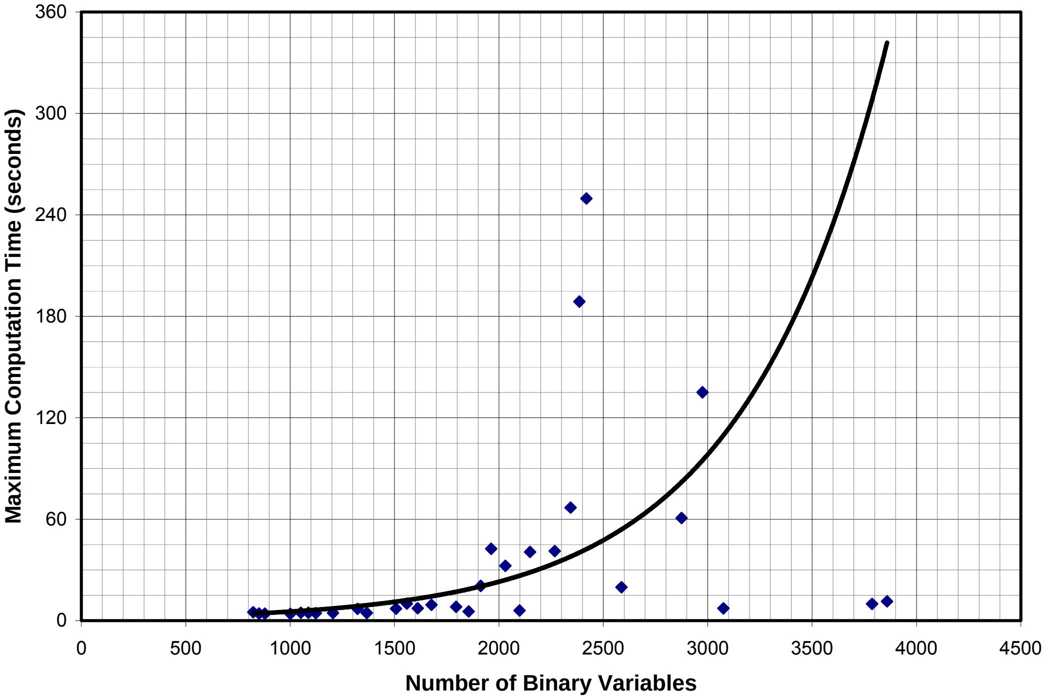

Figure 4.9. Maximum computation time of algorithm for various problem sizes. The solid black line is an exponential trend line.

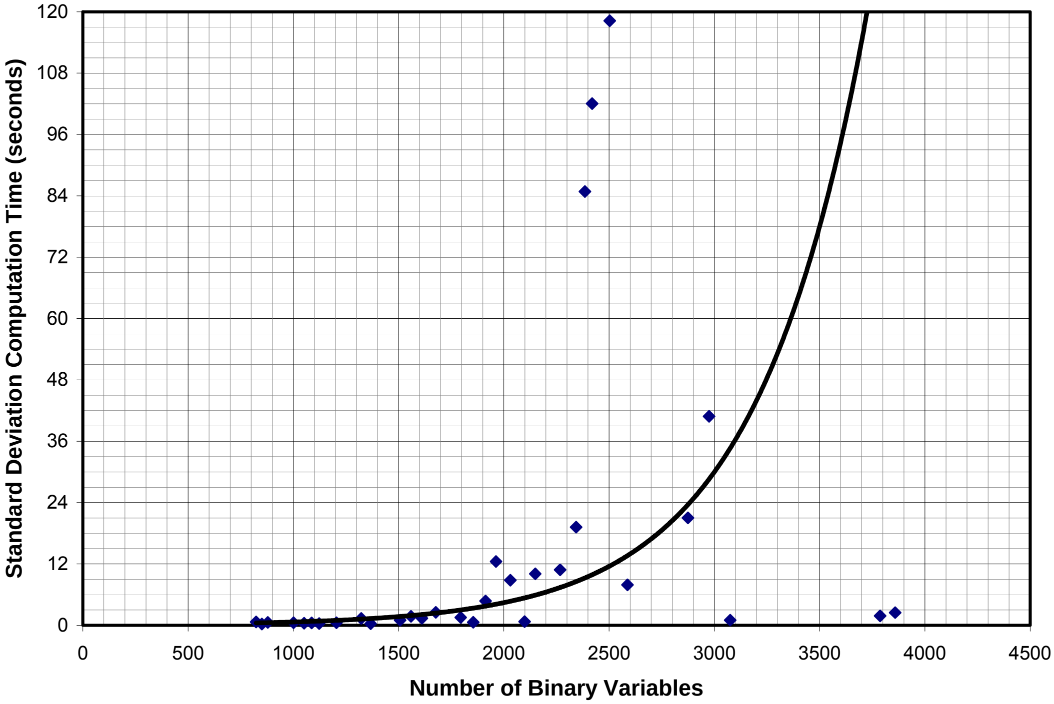

Figure 4.10. Standard deviation of computation time of algorithm for various problem sizes. The solid black line is an exponential trend line.

Chapter 5

Receding Horizon Control in the Presence of Multi-Scale Disturbances

5.1 Introduction

Receding horizon control, otherwise known as Model Predictive Control (MPC), has recently been studied as a controls technique for aerospace applications due to its ability to systematically handle vehicle dynamics constraints, performance limitations, and no-fly areas. Traditionally, MPC has been used in process control, where the problems are large and the time scales are typically on the order of minutes and hours [78]. Improvements in computer technology have enabled the use of receding horizon control in problems with much faster dynamics such as aircraft and spacecraft trajectory optimization [33].

As previously mentioned, the receding horizon controller divides the trajectory into three segments of varying fidelity to ensure the path planning problem studied by this research is tractable in real-time. The detailed trajectory is the segment extending from the initial position of the aircraft, and is computed subject to a vehicle dynamics and performance model. The number of waypoints in the detailed trajectory is called the planning horizon, and the waypoints are separated by a specified time step. The rest of the trajectory is represented by the cost-to-go, where the cost-to-go is a coarse approximation of the trajectory through a set of grid points extending away from the planning horizon. As described in Chapter 4, the cost-to-go represents the path of least cost from the end of the planning horizon to the destination. Figure 5.1 depicts the three segments of the trajectory, where the detailed trajectory is shown in green, the line-of-sight is in blue, and the cost-to-go is magenta. The trajectory is executed for a specified number of waypoints in the planning horizon, called the execution horizon, then the aircraft states and environment are updated, and the process is repeated until the destination is reached.

This chapter examines the effect of multi-scale disturbances in the environment on the performance of the receding horizon controller. The problem of multiple scales has been addressed in the process control industry as a function of time scale [79], but it has not been addressed in the context of a trajectory planning problem, which operates in a far more reactive fashion than was addressed previously. This research investigates the effect of multiple scales, specifically the effect of disturbance size and motion in a system with much faster dynamics than previously studied.

The scenarios of this chapter involve a single flight from Chicago to Los Angeles in the presence of obstacles that exploit the sensitivities of receding horizon trajectory planning. The disturbances studied in this chapter include hard and soft no-fly zone constraints which occur with varied areas and dynamics. The effect of disturbance size and speed on the trajectories will be initially analyzed independently, and then in a combined fashion. Lastly, a “best” receding horizon strategy is presented which provides a novel way to approach a multi-scale trajectory planning problem.

5.2 Sensitivity of Trajectory to Disturbance Size

Aircraft path planning problems can be subject to many different types of disturbances that are of diverse sizes. The path planning problem studied by this research specifically considers the disturbances of persistent contrail formation, thunderstorm avoidance, and collision avoidance. As mentioned in Chapter 2, persistent contrails form when an aircraft flies through an area in the atmosphere where the relative humidity with respect to ice is greater then 100%. These areas are extremely large in size, sometimes covering multiple states. Thunderstorms are much smaller, often with areas on the order of tens of square miles. Collision avoidance accounts for the smallest area, with avoidance regions of only a few square miles.

The pathological disturbances of this section are designed to emulate the characteristics of the events described above, and are generically classified as large, medium, small, and tiny. Table 5.1 shows the sizes of the obstacles used in this analysis, where the characteristic length is determined by the length of the widest section of the obstacle. Figure 5.2 depicts the definition of characteristic length for a typical obstacle, as is used in this chapter. This section will show how trajectories generated with receding horizon control are affected by obstacle size through varying parameters in the receding horizon controller and observing the results.

5.2.1 Effect of Planning Horizon Length

The planning horizon (N_p) is composed of the waypoints in the detailed trajectory section of the complete trajectory. The locations of the waypoints in the planning horizon are calculated subject to realistic vehicle dynamics and performance limitations. The length of the planning horizon is calculated by the following equation

$$l_p = \sum_{i=1}^{N_p} \left\| \vec{v}_i \right\| \cdot dt \tag{5.1}$$

where l_p is the planning horizon length, N_p is the number of waypoints in the planning horizon, $\boldsymbol{v}$ is the aircraft velocity, and dt is the time step. The independent variables in Eq. 5.1 are the number of waypoints in the planning horizon and the size of the time step. The sensitivity of these variables to the shape and fuel consumption of the trajectory will be analyzed.

5.2.1.1 Hard Avoidance Constraints

The first set of disturbances in the analysis includes obstacles with hard avoidance constraints, which enforce that no feasible trajectory can include a waypoint within the boundaries of the obstacle. Figure 5.3 shows the behavior of receding horizon trajectories with different planning horizons around a large static obstacle. The trajectories should be compared to the most optimal solution, which was generated as a non-receding horizon MILP and is depicted in blue. Additionally, it should be noted that the receding horizon trajectories were generated with an execution horizon of two waypoints and a time step of 2.4 minutes. The set of four trajectories in Fig. 5.3 take two distinct paths around the obstacle; the optimal and 24 waypoint planning horizon trajectories follow similar paths north of the obstacle, while the 12 and 6 waypoint planning horizon trajectories follow paths south of the obstacle. This is significant in that the 12 and 6 waypoint trajectories burn on average 3.0% more fuel than the optimal trajectory, compared to the 24 waypoint trajectory which burns 1.4% more fuel than the optimal trajectory. Note that the length of the 24 waypoint planning horizon is longer than the characteristic length of the obstacle.

Figure 5.4 shows the sensitivity of different planning horizons to a disturbance that is a medium sized obstacle with hard avoidance constraints. Similarly to Fig. 5.3, the trajectories in this set follow two distinct paths, where the trajectory with the highest amount of waypoints in the planning horizon more closely matches the optimal trajectory. The trajectory with 16 waypoints in the planning horizon only burned 0.8% more fuel than the optimal trajectory, while the trajectories with 8 and 12 waypoints in the planning horizon burned on average 2.6% more fuel than the optimal trajectory.

Figure 5.5 shows the paths of receding horizon trajectories with varied planning horizons around a small obstacle. In this example, all of the trajectories follow a similar path from the origin to the destination. There are occasional differences between the receding horizon trajectories and the optimal trajectory, but the differences are significantly less than in Figs. 5.3 and 5.4. This is most likely due to the fact that the characteristic length of the obstacle is much smaller than the lengths of the planning horizons.

The tiny class of obstacle presents an anomalous case that is designed to emulate collision avoidance in a multi-aircraft scenario. The characteristic length for the tiny obstacle is only 10 nautical miles, which corresponds to a “safety zone” of 5 nm radius around the aircraft. What makes this scenario unique is that if the time step is too large, the path planning algorithm can find a feasible trajectory through the safety zone by placing its waypoints on either side of the safety zone. To prevent this, the following relationship must be true:

$$\lambda > V_{\max} dt \tag{5.2}$$

where λ is the characteristic length, V_{max} is the maximum velocity, and *dt* is the size of the time step. If Eq. 5.2 is false, then the planning horizon can straddle the obstacle, as seen in Fig. 5.6 where the planning horizon waypoints are shown as orange stars.

5.2.1.2 Soft Avoidance Constraints

Some situations do not require the strict avoidance that hard constraints provide. In these cases, soft avoidance constraints are used that penalize obstacle penetration instead of precluding it. Figure 5.7 shows the behavior of receding horizon trajectories of different planning horizon length around a large obstacle with soft avoidance constraints. The receding horizon trajectories are compared to a non-receding horizon MILP trajectory, which is shown in blue. The trajectory with 8 waypoints in the planning horizon does not change course and passes through the obstacle, leaving 417 nm of path length within the boundaries of the obstacle. Because this trajectory does not alter its course around the obstacle, it burns 0.4% fuel less than the non-receding horizon trajectory. The 12 and 16 waypoint planning horizon trajectories climb above the obstacle, leaving only 74 and 63 nm of path length within the obstacle, respectively. However, these trajectories burn more fuel than the 8 waypoint trajectory; specifically 0.1% and 0.09% less fuel than the non-receding horizon trajectory, respectively.

Figure 5.8 shows the sensitivity of different planning horizons to a disturbance that is a medium sized obstacle with soft avoidance constraints. Similar to the large obstacle case, the trajectory with 8 waypoints in the planning horizon does not alter its course in response to the obstacle. This trajectory is responsible for 202 nm of path length within the obstacle and it burns 0.3% fuel less than the non-receding horizon trajectory. The trajectories with 12 and 16 waypoints in the planning horizon enter a slow climb to ascend about the obstacle but still create 115 and 111 nm of path length within the obstacle boundaries, respectively. These trajectories burn 0.04% and 0.03% less fuel than the non-receding horizon trajectory, respectively.

Figure 5.9 presents the effect of different planning horizons on the trajectory in the presence of a small sized obstacle with soft avoidance constraints. In this example none of the trajectories are affected by the presence of the obstacle. Planning horizon length has no appreciable effect on the overall shape and performance of the trajectories, as they all are responsible for approximately 65 nm of path length within the obstacle and fuel burn differences are negligible.

5.2.1.3 Summary of the Effect of Planning Horizon Length

The effect of planning horizon length is clearly evident in the examples presented in the prior subsection. Generally speaking, the longer the planning horizon, the better the receding horizon trajectory matches the optimal trajectory. Table 5.2 gives the maximum planning horizon length for planning horizons with different numbers of waypoints. The maximum planning horizon length was calculated with Eq. 5.1 with a maximum velocity of 470 knots and a time step of 2.4 minutes.

In the case of hard obstacle avoidance, it is interesting to note that as long as the maximum planning horizon length given in Table 5.2 is longer than the characteristic length of the obstacle given in Table 5.1, the receding horizon trajectory tracks the optimal trajectory relatively well. If the maximum length of the planning horizon is less than the characteristic length, then the receding horizon controller typically picks an inefficient path around the obstacle.

For the tiny obstacle case, in order for the inequality of Eq. 5.2 to be true, the time step cannot be greater than 1.28 minutes with a maximum velocity of 470 knots for the given aircraft models. Most of the analysis of this section was computed with a time step of 2.4 minutes, which is obviously too large to handle a tiny obstacle. If the time step of the earlier analysis were to change to 1.2 minutes, twice the number of waypoints would be required in the planning horizon to have the same maximum planning horizon length. For example, a trajectory with 16 waypoints in the planning horizon and a time step of 2.4 minutes was found to track the non-receding horizon trajectory well for medium sized and smaller obstacles. To get this same behavior with a time step of 1.2 minutes, the number of waypoints in the planning horizon would need to increase to 32. As mentioned in Chapter 4.4, the MILP method does not scale well with the number of binary variables, and increasing the number of waypoints introduces more binary variables to the optimization which could drive the algorithm to become intractable.

In the case of soft obstacle avoidance, there is a less obvious correlation between maximum planning horizon length and trajectory optimality. It is still generally true that the receding horizon trajectories with longer planning horizon lengths better tracked the non-receding horizon trajectory, but there is not an observable correlation between maximum planning horizon length, obstacle characteristic length, and overall trajectory shape.

5.2.2 Effect of Execution Horizon Length

The execution horizon (*N_e*) is composed of waypoints in the detailed trajectory that are visited before the receding horizon controller updates the environment. For example, if the execution horizon is two, the first two waypoints in the planning horizon are visited and then the trajectory is recalculated. The length of the execution horizon is a function of the following equation

$$l_e = \sum_{i=1}^{N_e} \left\| \vec{v}_i \right\| \cdot dt \tag{5.3}$$

where l_e is the execution horizon length, N_e is the number of waypoints in the execution horizon, $\boldsymbol{v}$ is the aircraft velocity, and *dt* is the time step. The independent variables in Eq. 5.3 are the number of waypoints in the execution horizon and the size of the time step. The sensitivity of these variables to the overall performance of the trajectory will be analyzed.

5.2.2.1 Hard Avoidance Constraints

The first set of disturbances in the analysis includes obstacles with hard avoidance constraints. Figure 5.10 shows the behavior of receding horizon trajectories with different numbers of waypoints in the execution horizon around a large obstacle with hard constraints. The planning horizon for this example consists of 12 waypoints and the time step is 2.4 minutes. The most obvious observation of this figure is that none of the receding horizon trajectories follow the path of the non-receding horizon trajectory, which is shown in blue. A closer look at the trajectories shows that the trajectories with the smaller execution horizons make an earlier decision to change course around the obstacle. In this case, the earlier decision corresponds to better performance. The trajectory with 2 waypoints in the execution horizon burned 2.8% more fuel than the non-receding horizon trajectory, compared to 3.0% and 3.1% for the 6 and 8 execution horizon trajectories, respectively.

Figure 5.11 presents the sensitivity to execution horizon for receding horizon trajectories around a medium sized obstacle with hard avoidance constraints. The planning horizon for this example consists of 16 waypoints and the time step is 2.4 minutes. All of the trajectories generally follow the same path, however there are small differences between the receding horizon trajectories and the non-receding horizon trajectory. The trajectory with 8 waypoints in the execution horizon follows a much more conservative path around the obstacle compared to the trajectories with 2 and 12 waypoints in the execution horizon. The trajectory with 2 waypoints in the execution horizon burned 0.8% more fuel than the non-receding horizon trajectory, while the trajectories with 8 and 12 waypoints in the execution horizon burned 1.3% and 0.7% more fuel than the non-receding horizon trajectory.

Figure 5.12 gives the sensitivity of the receding horizon trajectories to execution horizon in the presence of a small sized obstacle with hard avoidance constraints. The planning horizon for this example consists of 12 waypoints and the time step is 2.4 minutes. The trajectories in this figure are generally well behaved with the exception of the trajectory with 6 waypoints in the execution horizon, which alters its path to pass below the obstacle instead of above it. The trajectory with 8 waypoints in the execution horizon initially follows the 6 waypoint trajectory, but switches its course over north western Missouri to track closer to the non-receding horizon trajectory. The trajectory with 2 waypoints in the execution horizon is the closest to the non-receding horizon trajectory, and burns only 0.06% more fuel. The 6 waypoint trajectory was the most inaccurate, with a fuel burn of 1.6% more than the non-receding horizon trajectory. The 8 waypoint trajectory burned 0.07% more fuel than the non-receding horizon trajectory.

5.2.2.2 Soft Avoidance Constraints

The effect of varying execution horizon in the presence of soft avoidance constraints is discussed in this section. Figure 5.13 shows the sensitivity of receding horizon trajectories to execution horizon length in the presence of a large obstacle with soft avoidance constraints. In this example there are 16 waypoints in the planning horizon and the time step is 2.4 minutes. Generally speaking, the longer execution horizons correspond to trajectories closer to the non-receding horizon trajectory. The trajectory with 2 waypoints in the execution horizon consists of 63 nm of path length within the boundaries of the obstacle, whereas the trajectories with 8 and 12 waypoints in the execution horizon correspond to approximately 8 and 6 miles of path length in the obstacle, respectively. The 2 waypoint trajectory burned 0.06% less fuel than the non-receding horizon trajectory, the difference in fuel burn between the 8 waypoint trajectory and the non-receding horizon trajectory is negligible, and the 12 waypoint trajectory burned 0.05% less fuel than the non-receding horizon trajectory.

Figure 5.14 shows the sensitivity of receding horizon trajectories to execution horizon length in the presence of a medium obstacle with soft avoidance constraints. In this example there are 16 waypoints in the planning horizon and the time step is 2.4 minutes. The trajectory with 2 waypoints in the execution horizon created a path length of 115 nm and burned 0.3% less fuel than the non-receding horizon trajectory. The trajectories with 6 and 8 waypoints in the execution horizon created path lengths of 8 nm and 43 nm, respectively. The 6 waypoint trajectory burned 0.02% more fuel than the non-receding horizon trajectory, and the 8 waypoint trajectory burned 0.06% less fuel than the non-receding horizon trajectory.

Figure 5.15 presents the behavior of receding horizon trajectories of different execution horizon lengths around a small sized obstacle with soft avoidance constraints. The receding horizon trajectories in this figure have 12 waypoints in the planning horizon and use a time step of 2.4 minutes. All of the trajectories penetrate the obstacle without altering their course and create 65 nm of path length within the boundaries of the obstacle. The fuel burn of all of the trajectories exhibit negligible differences.

5.2.2.3 Summary of the Effect of Execution Horizon Length

Simulations were conducted to analyze the sensitivity of receding horizon trajectories to the length of the execution horizon (Eq. 5.3). The effect of execution horizon length on the receding horizon trajectories is more subtle than the effect of planning horizon length; however, a definite trend exists. In general, the trajectories with smaller execution horizons performed more efficiently than those with longer execution horizons when in the presence of hard avoidance constraints. When navigating through areas with soft avoidance constraints, the opposite is true; the trajectories with longer execution horizons more closely track the non-receding horizon trajectory.

In the case of hard avoidance constraints, the shorter execution horizons were more efficient than the longer execution horizons. This is explained by first realizing that the receding horizon controller updates the trajectory at the end of each execution horizon, which implies that the trajectories with the shorter horizons are updated more frequently than the longer horizons. More frequent updates to the environment allow the aircraft to potentially make an earlier decision to divert around the obstacle. However, this is not an absolute rule, as seen in Fig. 5.11, where the trajectory with the longest execution horizon tracks the non-receding horizon trajectory as well as the trajectory with the shortest execution horizon. The conclusion should be that the trajectories with shorter execution horizon are more *probable* to make the most efficient decision to navigate around an obstacle with hard avoidance constraints.

In the case of soft avoidance constraints, the longer execution horizon trajectories are observed to be closer to the non-receding horizon trajectory. There are possible explanations for this behavior. First, the typical avoidance or mitigation maneuver for these types of obstacles is a climb or descent due to the fact that the obstacles are usually thin vertically and wide horizontally (they emulate areas of ice supersaturation). It is possible that the receding horizon strategy is not as effective in the vertical plane compared to the horizontal plane due to the extremely disparate scales of the vertical and horizontal motion of aircraft. Second, the acceleration penalty in the detailed trajectory optimization might slow the climb and descent rates in the first few waypoints of the planning horizon. If this is the case, the trajectories with short execution horizons cannot build as high of a rate of climb as the trajectories with long execution horizons, leading to the behavior seen in Figs. 5.13 and 5.14.

5.3 Sensitivity of Trajectory to Disturbance Dynamics

Until now, the disturbances analyzed in this chapter have been static. Dynamic obstacles add an additional dimension of realism and difficulty to the path planning problem. This section will investigate the effect of planning horizon length and execution horizon length to obstacles that are dynamic. The pathological obstacles of this section were designed to emulate the size and movement of thunderstorm cells, and are generically classified as large, medium, and small. As before, Table 5.1 shows the sizes of the obstacles used in this analysis, where the characteristic length is determined by the length of the widest section of the obstacle.

5.3.1 Effect of Planning Horizon Length

The effect of planning horizon length on the overall shape and performance of receding horizon trajectories around dynamic obstacles of different scales is analyzed in this section. The locations of the waypoints in the planning horizon are calculated subject to realistic vehicle dynamics and performance limitations. See Chapter 4 for a detailed description of how the receding horizon controller handles dynamic obstacles.

Figure 5.16 shows the behavior of receding horizon trajectories of different planning horizon length in the presence of a dynamic large sized obstacle with hard avoidance constraints. The avoidance region is the red box and there is a map of the United

States in the background to provide perspective. The execution horizon for this example consists of 2 waypoints and the time step is 2.4 minutes. The obstacle is moving in a south-eastern direction. The non-receding horizon trajectory is theoretically the most optimal trajectory in the figure and is shown as the blue dashed line. The trajectory with 24 waypoints in the planning horizon follows most closely to the non-receding horizon trajectory, which passes to the north of the obstacle. The trajectories with 8 and 16 waypoints in the planning horizon pass to the south of the obstacle, which is in the obstacle's direction of motion. The southern trajectories travel a much farther distance to get to the destination, and this corresponds to 10.1% and 9.8% more fuel burned than in the non-receding horizon trajectory. In addition those trajectories also arrive approximately 18 and 15 minutes after the non-receding horizon trajectory. The trajectory with 24 waypoints in the planning horizon only burned 2.6% more fuel than the non-receding horizon trajectory, and arrived 4 minutes later.

Figure 5.17 gives the sensitivity of receding horizon trajectories to different planning horizon lengths in the presence of a dynamic medium sized obstacle with hard avoidance constraints. The avoidance region is the red box and there is a map of the United States in the background to provide perspective. The execution horizon for this example consists of 2 waypoints and the time step is 2.4 minutes. As before, the non-receding horizon trajectory is shown as the blue dashed line. The trajectories with 16 and 12 waypoints in the planning horizon follow closely to the non-receding horizon trajectory, although the 16 waypoint trajectory is the closest. The shorter planning horizon lengths, specifically 6 and 8 waypoints in this case, choose a path to the south of the obstacle. The longer paths of the 6 and 8 waypoint trajectories correspond to

12.1% and 12.3% more fuel burned when compared to the non-receding horizon trajectory, respectively. In addition those trajectories arrive approximately 24 and 25 minutes after the non-receding horizon trajectory. The trajectories with 12 and 16 waypoints in the planning horizon burned 2.8% and 1.7% more fuel than the non-receding horizon trajectory, respectively. Also, when compared to the non-receding horizon trajectory, the 12 and 16 waypoint trajectories arrive 6 and 8 minutes later.

Figure 5.18 presents the effect of different planning horizon lengths on receding horizon trajectories in the presence of a dynamic small sized obstacle with hard avoidance constraints. The avoidance region is the red box and there is a map of the United States in the background to provide perspective. The execution horizon for this example consists of 2 waypoints and the time step is 2.4 minutes. In this example, the trajectories with 8, 12, and 16 waypoints in the planning horizon all follow the same general path as the non-receding horizon trajectory. Conversely, the trajectory with 6 waypoints in the planning horizon passes to the south of the obstacle. The southern deviation is not as severe in this example compared to the others because the obstacle is smaller. The 6 waypoint trajectory burns 3.2% more fuel than the non-receding horizon trajectory. The 8, 12, and 16 waypoint trajectories burn a negligible amount compared to the non-receding horizon trajectory, and all trajectories arrive at the destination at approximately the same time.

5.3.2 Effect of Execution Horizon Length

The effect of execution horizon length on the overall shape and performance of receding horizon trajectories around dynamic obstacles of different scales is analyzed in this section. As previously mentioned, the execution horizon is the number of waypoints in the planning horizon that are visited before the receding horizon controller updates the trajectory. The length of the execution horizon was defined earlier in Eq. 5.2, and is dependent on the number of waypoints, the speed of the aircraft, and the size of the time step.

Figure 5.19 shows the behavior of receding horizon trajectories with different execution horizons around a large dynamic obstacle with hard avoidance constraints. The avoidance region is the red box and there is a map of the United States in the background to provide perspective. The planning horizon for this example consists of 12 waypoints and the time step is 2.4 minutes. In this example none of the receding horizon trajectories follow the general shape of the non-receding horizon trajectory. Instead of changing course to the north to avoid the obstacle, they divert to the south, which is in the path of motion of the obstacle. This significantly affects the trajectories where the 2, 6, and 8 waypoint trajectories burned 9.8%, 9.9%, and 9.9% more fuel than the non-receding horizon trajectory, respectively. In addition, the receding horizon trajectories arrived at the destination 17, 18, and 18 minutes after the non-receding horizon trajectory, respectively.

Figure 5.20 gives the sensitivity of receding horizon trajectory performance to different execution horizons in the presence of a medium dynamic obstacle with hard avoidance constraints. The avoidance region is the red box and there is a map of the United States in the background to provide perspective. The planning horizon for this example consists of 12 waypoints and the time step is 2.4 minutes. The receding horizon trajectories with longer execution horizons do not reroute south of the obstacle as they did in Fig. 5.19. Instead, they first adjust their route to go south of the obstacle, but then change direction close to the obstacle and pass to the north. The trajectory with an execution horizon of 2 waypoints was the closest to the non-receding horizon trajectory, burning only 1.8% more fuel and arriving 6 minutes later. The 3, 6, and 8 waypoint trajectories were not quite as close, and they burned 4.2%, 4.4%, and 4.5% more fuel than the non-receding horizon trajectory. Furthermore, the 3, 6, and 8 waypoint trajectories arrived 9, 10, and 10 minutes later than the non-receding horizon trajectory.

Figure 5.21 presents the effect of different execution horizon lengths on receding horizon trajectory performance in the presence of a small dynamic obstacle with hard avoidance constraints. The avoidance region is the red box and there is a map of the United States in the background to provide perspective. The planning horizon for this example consists of 12 waypoints and the time step is 2.4 minutes. In this example all of the receding horizon trajectories follow the same general path. The trajectory with 8 waypoints execution horizon shows a slight deviation over northeastern Kansas, and burned 1.1% more fuel than the non-receding horizon trajectory. The receding horizon trajectories with 2, 4, and 6 waypoints in the execution horizon burned approximately the same amount of fuel as the non-receding horizon trajectory and arrived approximately 2 minutes later.

An additional comment on the effect of execution horizon on receding horizon trajectories in a dynamic environment is that the feasibility of the trajectory can be called into question if the execution horizon is too long. If the environment is dynamic and unpredictable, a long execution horizon can make the trajectory vulnerable to obstacle penetration due to infrequent updates of the planning horizon. In other words, an obstacle could change its shape or dynamics and encroach on the aircraft before the receding horizon controller updates the environment. This behavior was observed during extensive simulation of test case scenarios performed during this chapter.

5.3.3 Summary of the Effect of Dynamic Disturbances

Simulations were performed to analyze the sensitivity of receding horizon trajectory shape (length and curvature) to planning and execution horizon length in an environment containing dynamic obstacles. This analysis only considered the case of dynamic hard avoidance constraints. The general sensitivity between trajectory shape and planning and execution horizon length in a dynamic environment is similar to the sensitivity in a static environment: the receding horizon trajectory becomes is more efficient with long planning horizons and short execution horizons. Similarly to the earlier analysis, the time step used in the simulations is 2.4 minutes and the maximum velocity of the aircraft is 470 knots.

The effect of planning horizon length on the trajectory shape was similar in trend but more substantial when compared to the static obstacle analysis. The trajectories with longer planning horizons than the characteristic length of the obstacle made better course adjustment decisions, and the trajectories with the shorter planning horizons typically made poor decisions. However, the poor decisions in the dynamic obstacle examples proved to be much more inefficient than in the static examples due to the fact that the poor decisions taken during the planning process often guided the aircraft in the path of the obstacle, further pushing the trajectory away from the optimal trajectory.

The effect of execution horizon length on the receding horizon trajectories was also similar to the static obstacle examples, but much more severe. When considering dynamic obstacles, especially ones where the motion might be unpredictable, it is extremely important to frequently update the environment in the receding horizon controller to maintain the feasibility of the trajectory. The update rate of the receding horizon controller is directly related to the size of the execution horizon and the size of the time step. Therefore, it is relatively straight forward to hypothesize that a receding horizon trajectory with a shorter execution horizon would perform better than a trajectory with a long execution horizon. This behavior was observed in the simulations, where in general an execution horizon greater than 2 waypoints (with a time step of 2.4 minutes) caused significant differences in the trajectory. In some cases the trajectory would be rendered infeasible due to the fact that the receding horizon controller would not update sufficiently quickly to track the progress of the dynamic obstacles.

5.4 Receding Horizon Strategy for Multi-Scale Disturbances

This chapter has shown that receding horizon control is very sensitive to planning horizon length, execution horizon length, and to the spatial and temporal scales of the environment. Therefore, it is hard to choose a single “best” receding horizon strategy for a path planning problem of multiple scales. The safest strategy would be to have a small time step and small execution horizon to ensure collision avoidance and trajectory feasibility in the presence of dynamic obstacles. However, it was shown earlier that a long planning horizon length is needed to have an efficient trajectory. The combination of a long planning horizon length and a small time step leads to many waypoints in the planning horizon, which is not efficient computationally (see Chapter 4.4), and potentially not tractable in a real time setting. Therefore, this research proposes an adaptive receding horizon strategy that switches between a set of planning horizons, execution horizons, and time step sizes based on the obstacle environment to ensure trajectory feasibility and to promote efficiency.

This strategy is based on the premise that the aircraft is aware of what type of obstacle it is closest to or most effected by, so that it can dynamically select the best planning horizon, execution horizon, and time step to fly the most efficient and safe trajectory. This strategy uses the results presented earlier in this chapter.

It was found that in the case of large and medium sized obstacles with soft avoidance constraints the most optimal trajectory is generated with a long planning horizon and a long execution horizon. This type of obstacle was designed to emulate an area of ice supersaturation, which is where persistent contrails are formed. These areas are the largest, and most consistent disturbance appearing in this path planning problem, therefore it will serve as the default scale, known in this dissertation as Mode (1). A second reason for making this the default mode is that persistent contrail formation is not a safety of flight issue, and therefore switching from its mode of operation does not introduce any risk to the trajectory.

The second scale considered by the controller is the case of a large, medium, or small obstacle with hard avoidance constraints. These obstacles are designed to emulate thunderstorms, which are more threatening to safety, and therefore this scale takes priority over the default scale. If the aircraft is within a specified radius of a thunderstorm, the receding horizon controller switches to a smaller execution horizon, which updates the environment more frequently and promotes a more efficient and safe trajectory around the obstacle. This scale is known as Mode (2).

The third and final scale considered by the controller is the case of a tiny obstacle with hard avoidance constraints. Avoidance of this obstacle is intended to imitate collision avoidance, which is the most safety critical, and therefore takes priority over all other scales. If the aircraft is within a specified radius of another aircraft, the receding horizon controller switches to a smaller time step, which updates the environment more frequently and promotes a safe route around the obstacle. This scale is known as Mode (3).

Simulations were performed to test the viability of this strategy. Table 5.3 presents the modes used by the receding horizon controller. Figure 5.22 shows the performance of the new receding horizon strategy compared to two examples with fixed planning horizon, execution horizon, and time step. The hard avoidance regions are shown as red boxes and the soft avoidance regions are shown as yellow boxes. The top subplot is a horizontal representation of the trajectories, shown with a map of the United States in the background. The bottom subplot is a vertical representation of the trajectories. The benefit of having a receding horizon strategy that can adapt to the type of obstacle in the environment is apparent in the figure, as the adaptive strategy shares the favorable characteristics of both the Mode (1) and Mode (2) trajectories.

The Mode (1) trajectory was generated with a planning horizon of 12 waypoints and an execution horizon of 6 waypoints. This trajectory is designed with receding horizon characteristics that perform well with soft avoidance areas, and this behavior is seen by its ability to climb over the last soft obstacle with a minimal path length in the obstacle.

However, the execution horizon of this trajectory is poorly suited for an environment with dynamic hard constraints, which is visible by the overly conservative route generated around the three red boxes.

The Mode (2) trajectory was optimized with a planning horizon of 12 waypoints and an execution horizon of 2 waypoints. This trajectory is suited to traverse an environment containing hard and dynamic obstacles because the receding horizon controller is able to update the environment frequently. The shape of this trajectory was as expected; it was much more efficient navigating the three red boxes than the Mode (1) trajectory. However, the Mode (2) trajectory was not as efficient with the soft obstacles as the Mode (1) trajectory. The Mode (2) trajectory consisted of 157 nm of path length in the last soft obstacle, compared to only 9 nm for the Mode (1) trajectory.

The adaptive strategy presented in this section was effective in combining the positive qualities of the two fixed parameter trajectories. The adaptive trajectory started with the same planning horizon and execution horizon as the Mode (1) trajectory, and then switched to the Mode (2) planning and execution horizon when the hard obstacles were within 600 nm of the aircraft. At this point, the trajectory changes from the Mode (1) trajectory to the Mode (2) trajectory, which can be seen in Fig. 5.22 at a position over the western border of Missouri. The adaptive trajectory maintains the Mode (2) planning and execution horizons until the aircraft has passed the red boxes, at which point the planning and execution horizon switch back to the Mode (1) values. When the adaptive trajectory encounters the last soft obstacle, it behaves like the Mode (1) trajectory and climbs over the region with minimal path length in the obstacle.

The adaptive trajectory outperformed the Mode (1) or Mode (2) trajectories. It only burned 0.4% more fuel than the Mode (2) trajectory, and 3.6% less fuel than the Mode (1) trajectory. In addition, the adaptive trajectory created 17 nm of path length within the obstacles, compared to 202 nm for the Mode (2) trajectory and 9 nm for the Mode (1) trajectory, although it should be noted that the Mode (1) trajectory completely diverted around the first soft obstacle.

5.5 Summary

This chapter presented the sensitivity of the receding horizon control algorithm developed in this research to disturbances of multiple spatial and temporal scales. It was observed that, in general, long planning horizons perform better than short planning horizons. In addition, short execution horizons are superior when dynamic obstacles are present in the environment due to the faster update rate of the environment in the receding horizon controller. When considering the case of three-dimensional soft avoidance, longer execution horizons were observed to be more efficient in climbs and descents around the obstacles. An adaptive receding horizon strategy was developed from these results in an attempt to make a more general receding horizon controller, and the adaptive receding horizon strategy was shown to be both agile and effective in planning a path around obstacles of multiple scales.

Table 5.1. Characteristic lengths in nautical miles of the obstacles analyzed for this section.

	Soft (nm)	Hard (nm)
large	509.1	339.4
medium	349.9	226.8
small	189.7	108.2
tiny	N/A	10.0

Table 5.2. Maximum planning horizon lengths for number of waypoints in the planning horizon. Δt = 2.4 minutes.

N_p	Max Length (nm)
24	451.2
16	300.8
12	225.6
8	150.4
6	112.8
4	75.2

Table 5.3. A description of the modes used to switch between receding horizon parameters based on the scales in the environment.

Mode	(1)	(2)	(3)
N_p	12	12	12
N_e	6	2	2
Δt (min)	2.4	2.4	0.6

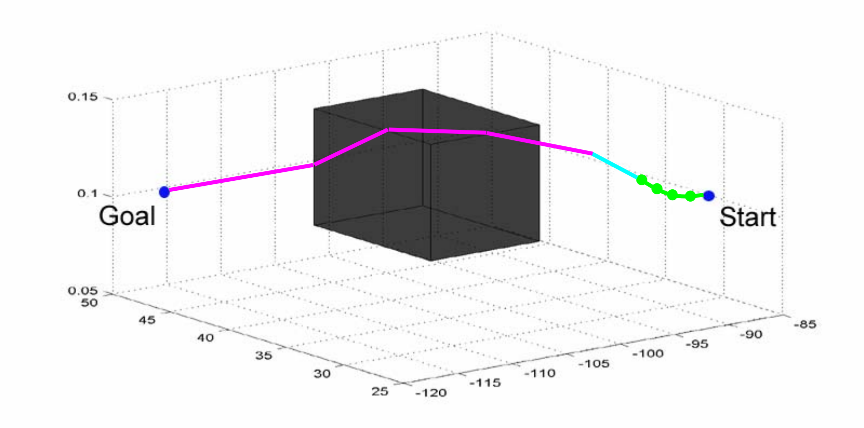

Figure 5.1. Simple example showing the relationship between the detailed trajectory (green), the line of sight (blue), and the cost-to-go (magenta).

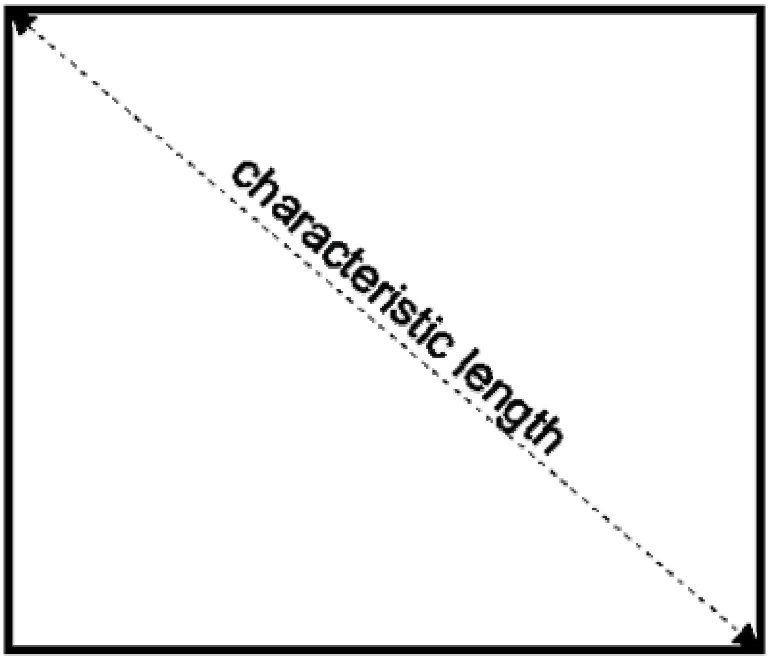

Figure 5.2. Depiction of the characteristic length for a generic square obstacle.

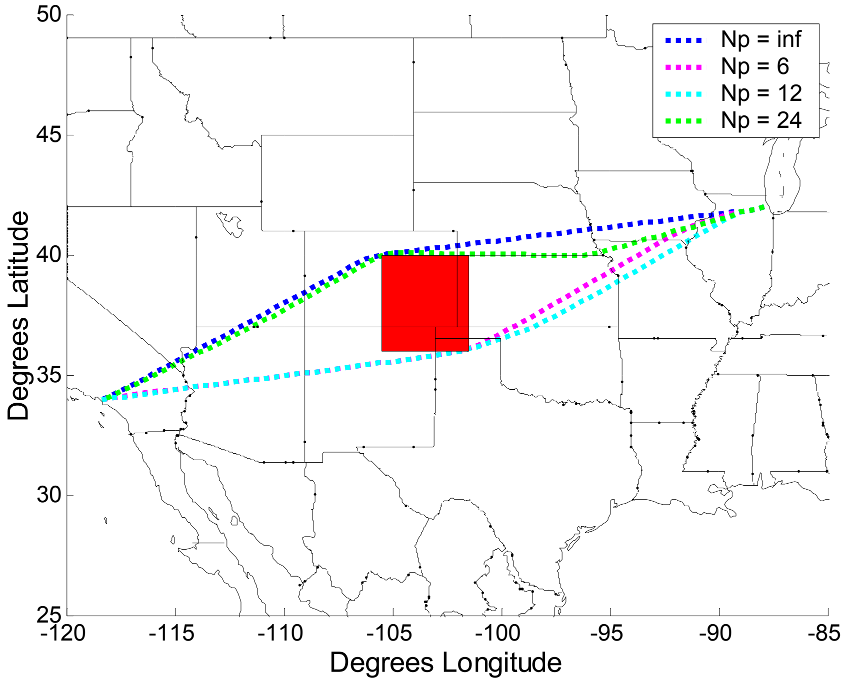

Figure 5.3. Comparison of trajectories (N_e = 2, dt = 2.4 min) around a large sized obstacle with hard avoidance constraints.

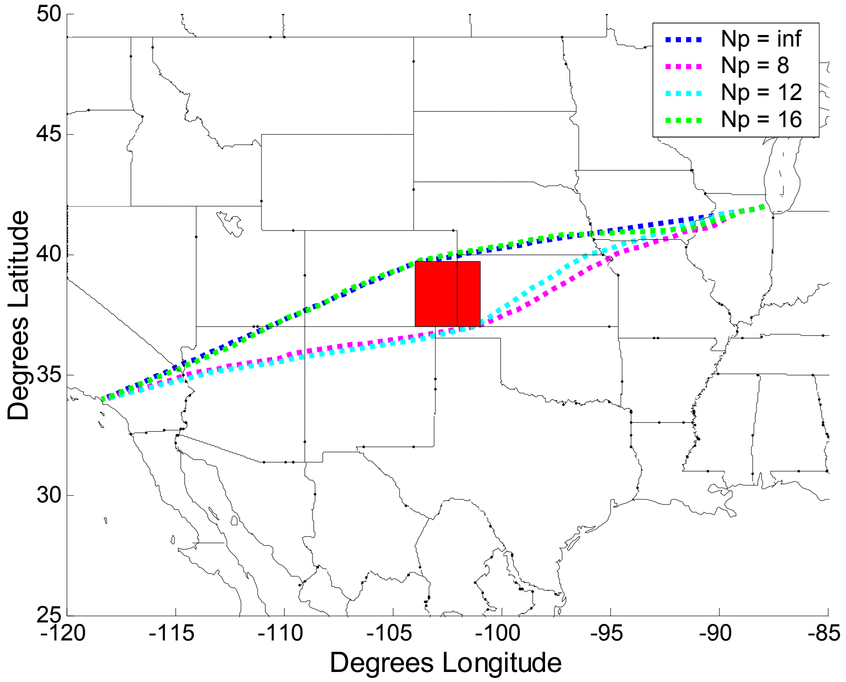

Figure 5.4. Comparison of trajectories (N_e = 2, dt = 2.4 min) around a medium sized obstacle with hard avoidance constraints.

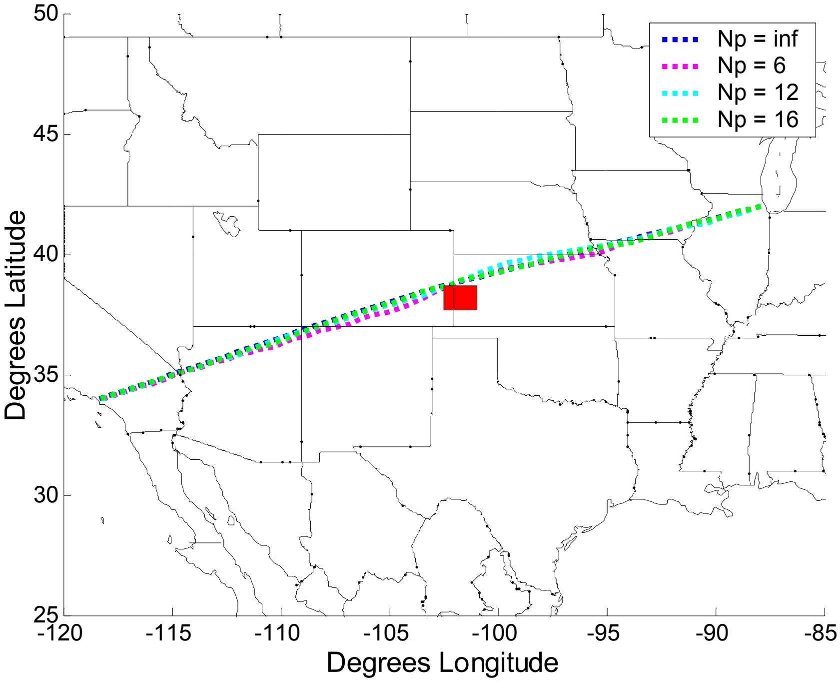

Figure 5.5. Comparison of trajectories (N_e = 2, *dt* = 2.4 min) around a small sized obstacle with hard avoidance constraints.

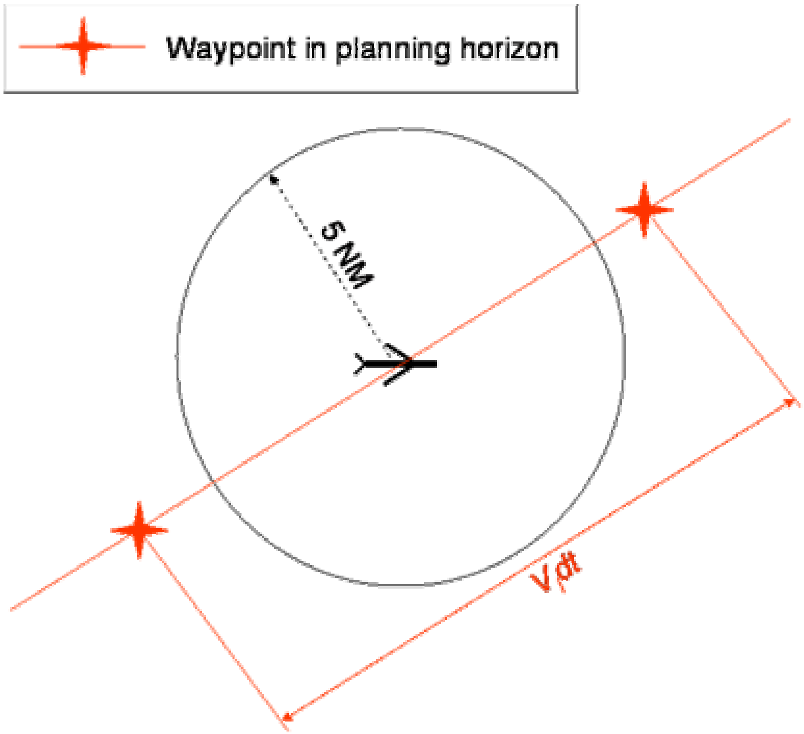

Figure 5.6. Situation in which the path planning algorithm finds a feasible trajectory through the safety zone of another aircraft because the time step is too large.

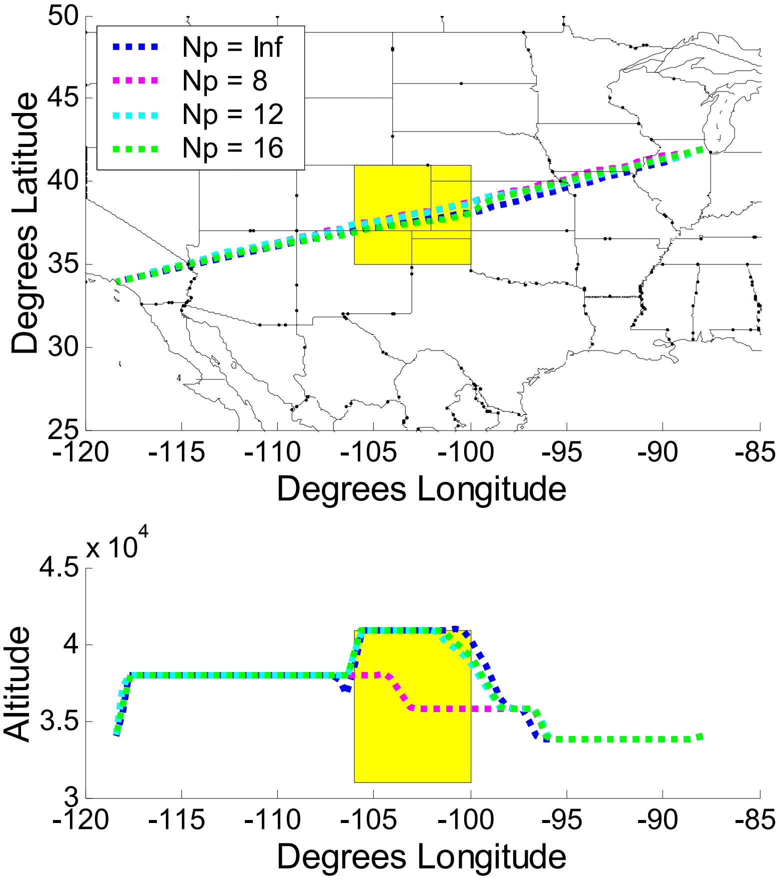

Figure 5.7. Comparison of trajectories (N_e = 2, dt = 2.4 min) around a large sized obstacle with soft avoidance constraints.

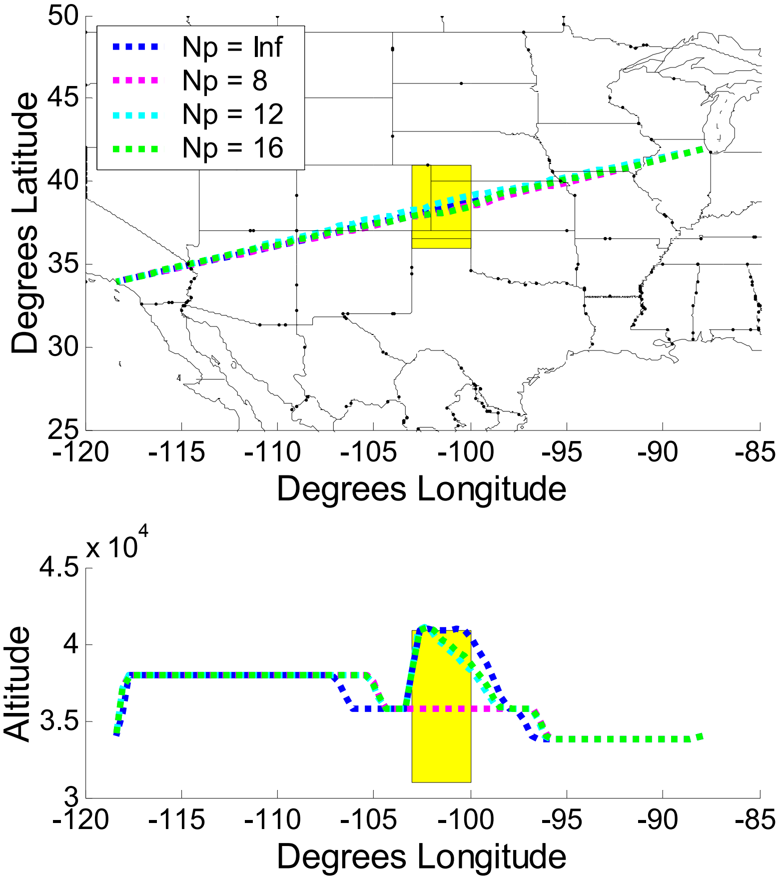

Figure 5.8. Comparison of trajectories (N_e = 2, dt = 2.4 min) around a medium sized obstacle with soft avoidance constraints.

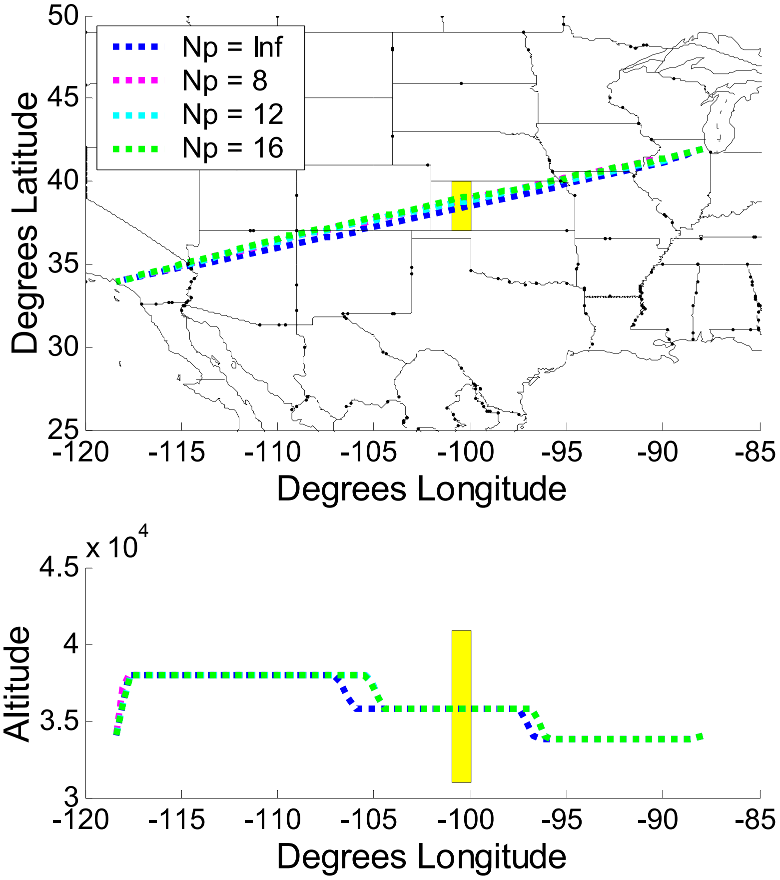

Figure 5.9. Comparison of trajectories (N_e = 2, dt = 2.4 min) around a small sized obstacle with soft avoidance constraints.

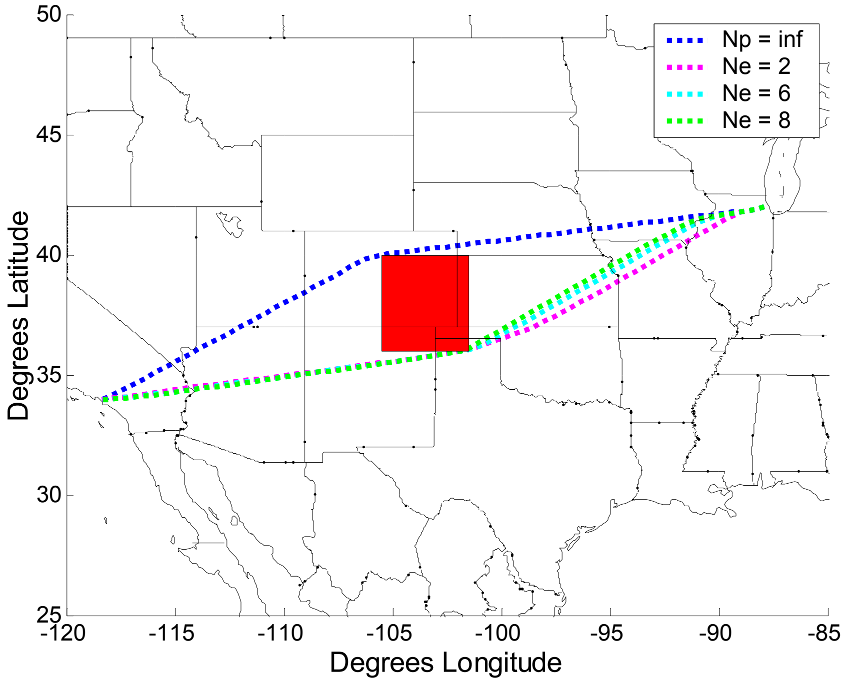

Figure 5.10. Comparison of trajectories (N_p = 12, dt = 2.4 min) around a large sized obstacle with hard avoidance constraints.

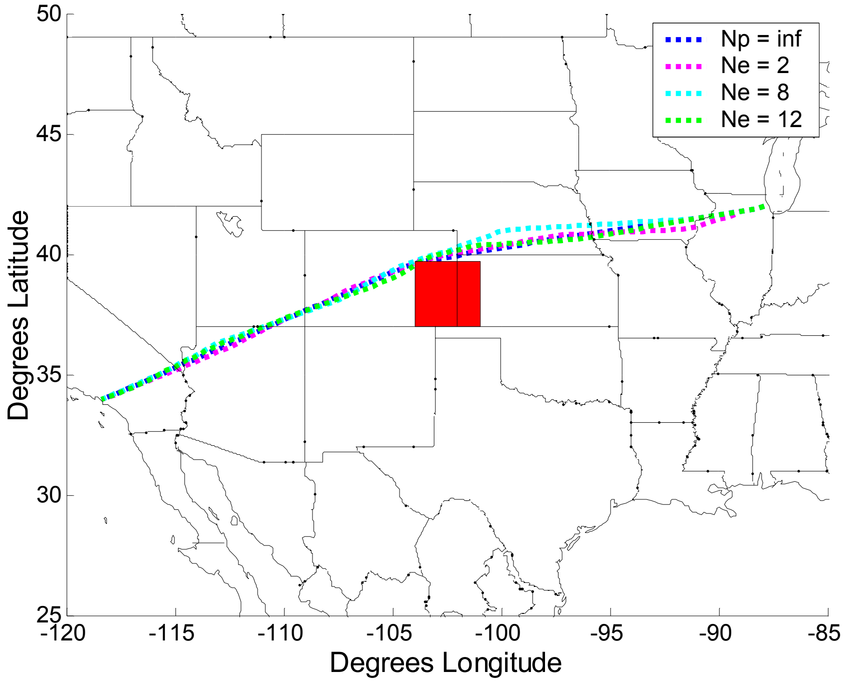

Figure 5.11. Comparison of trajectories (N_p = 16, dt = 2.4 min) around a medium sized obstacle with hard avoidance constraints.

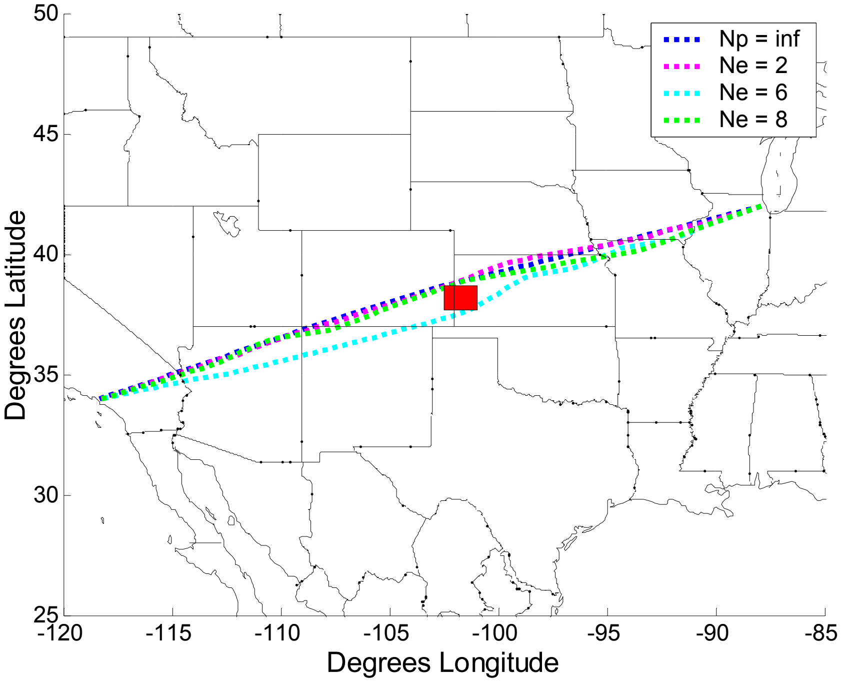

Figure 5.12. Comparison of trajectories (N_p = 12, dt = 2.4 min) around a small sized obstacle with hard avoidance constraints.

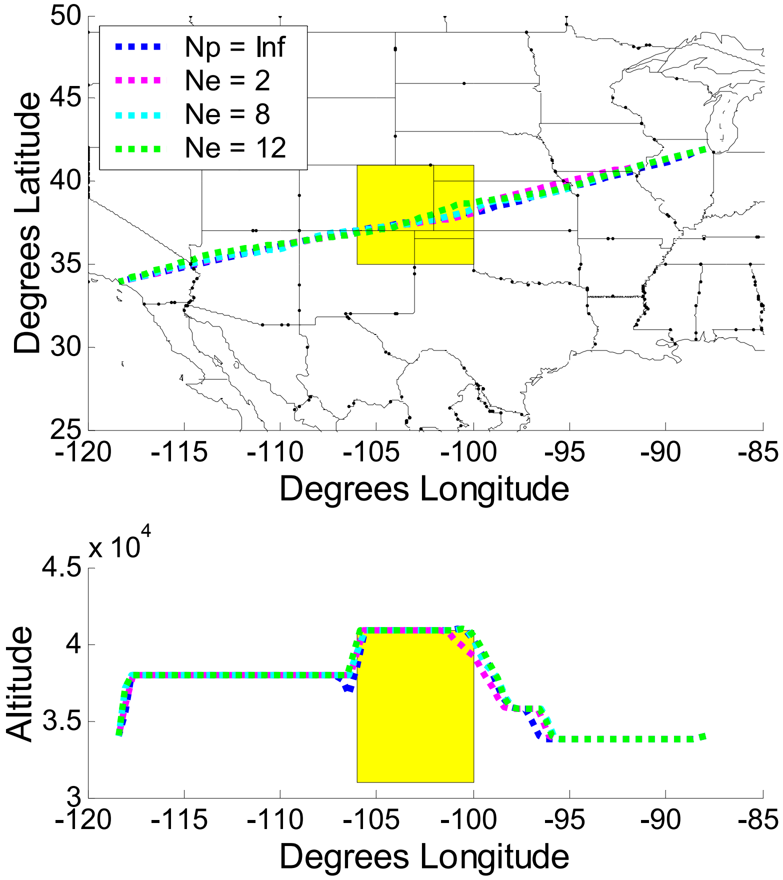

Figure 5.13. Comparison of trajectories (N_p = 16, dt = 2.4 min) around a large sized obstacle with soft avoidance constraints.

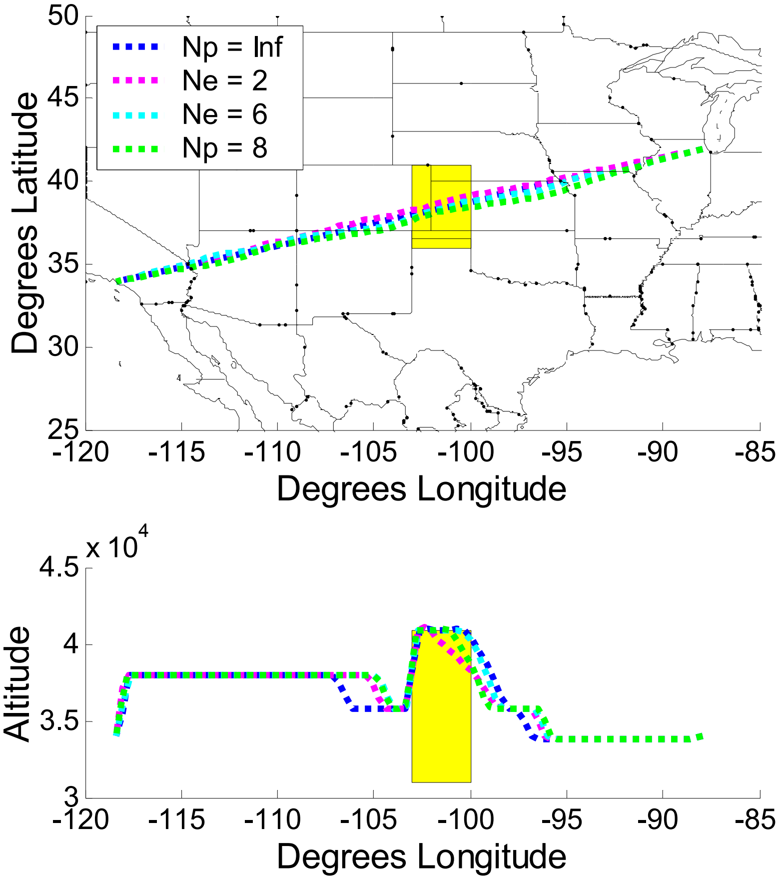

Figure 5.14. Comparison of trajectories (N_p = 12, *dt* = 2.4 min) around a medium sized obstacle with soft avoidance constraints.

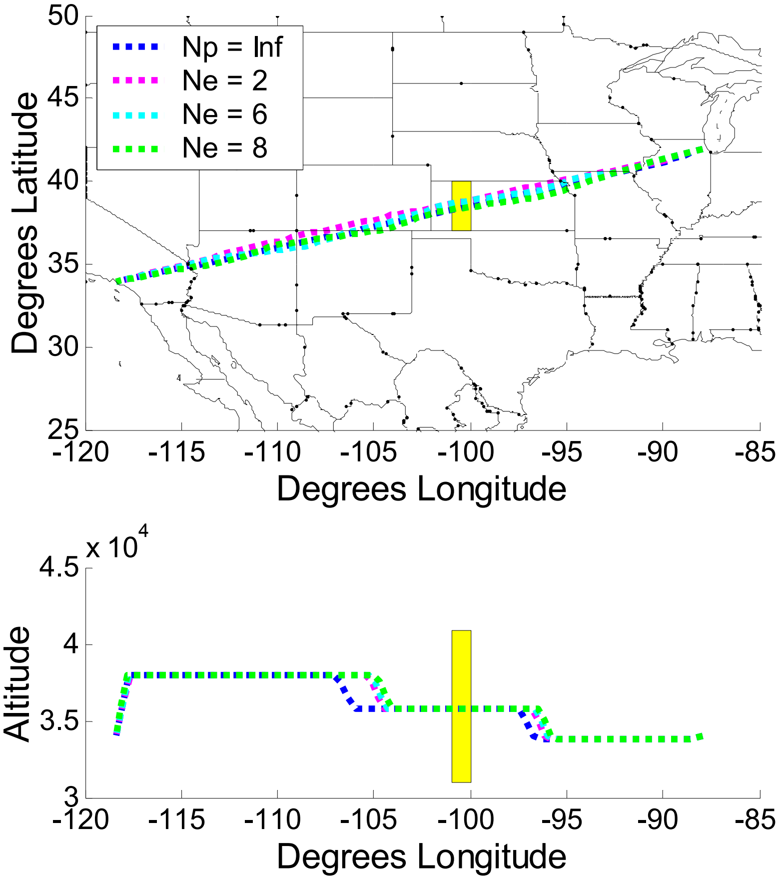

Figure 5.15. Comparison of trajectories (N_p = 12, *dt* = 2.4 min) around a small sized obstacle with soft avoidance constraints.

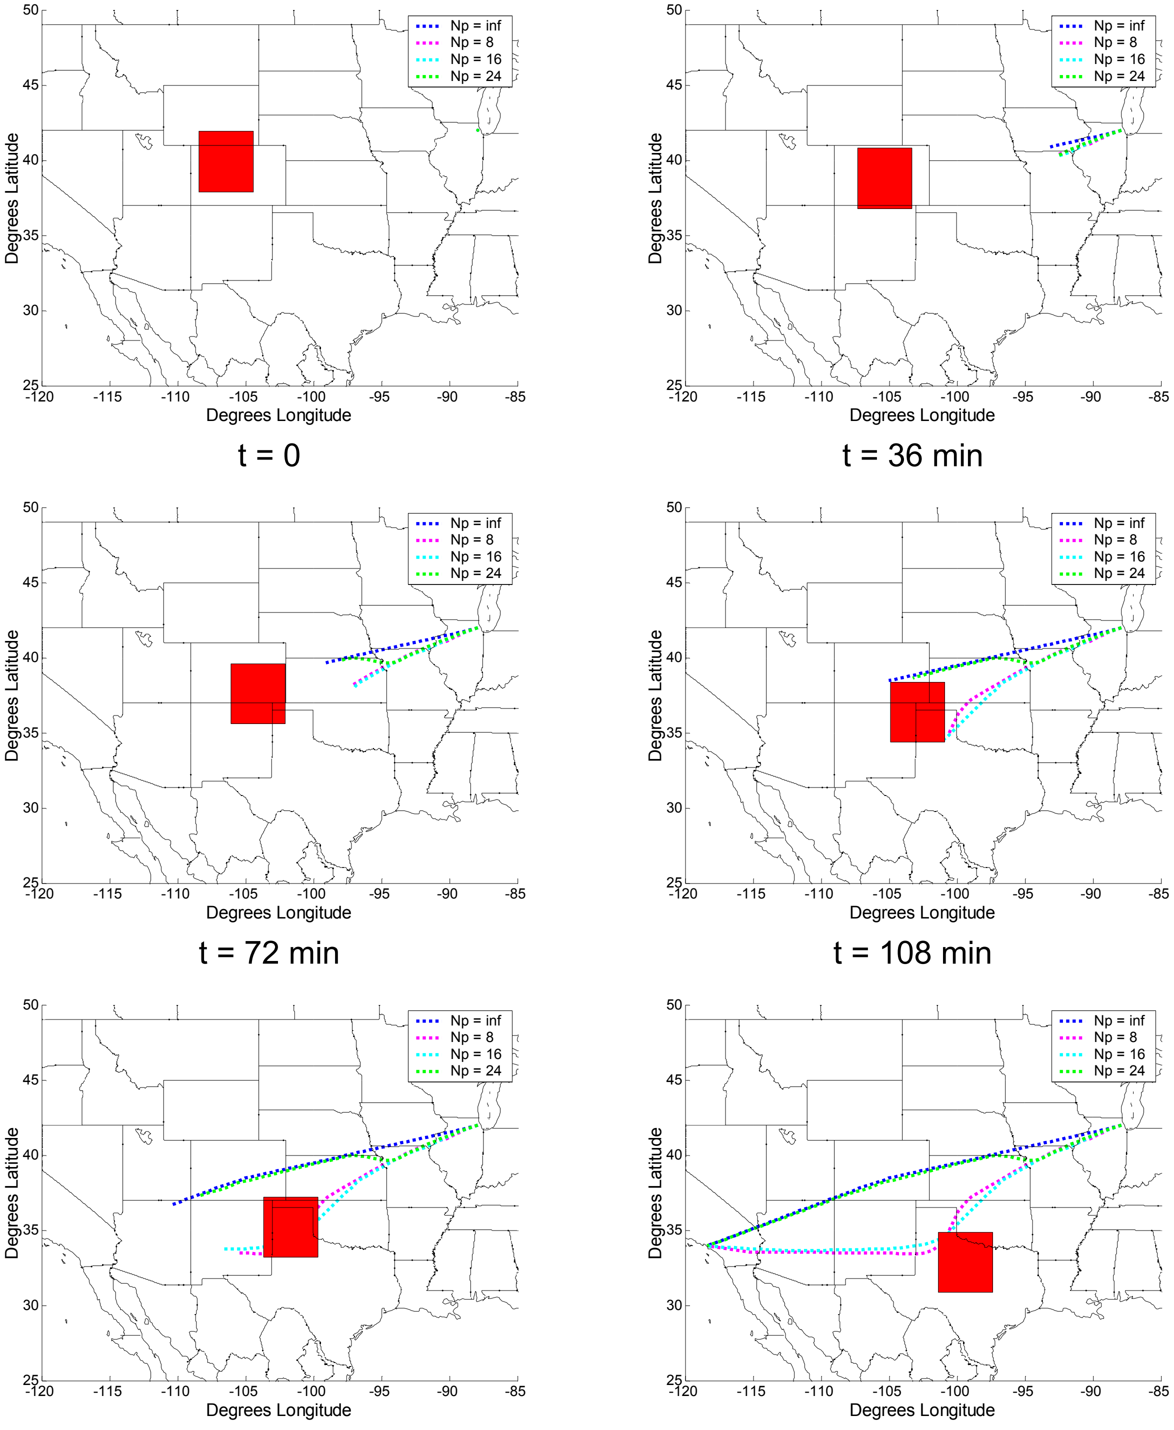

t = 0

t = 36 min

t = 72 min

t = 108 min

t = 144 min

t = end

Figure 5.16. Comparison of trajectories (N_e = 2, dt = 2.4 min) around a large sized dynamic obstacle with hard avoidance constraints.

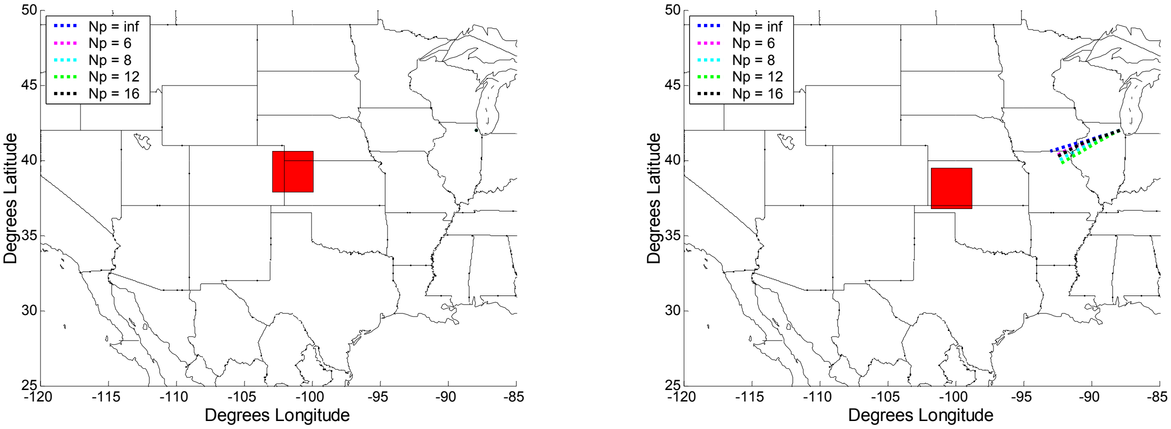

t = 0

t = 36 min

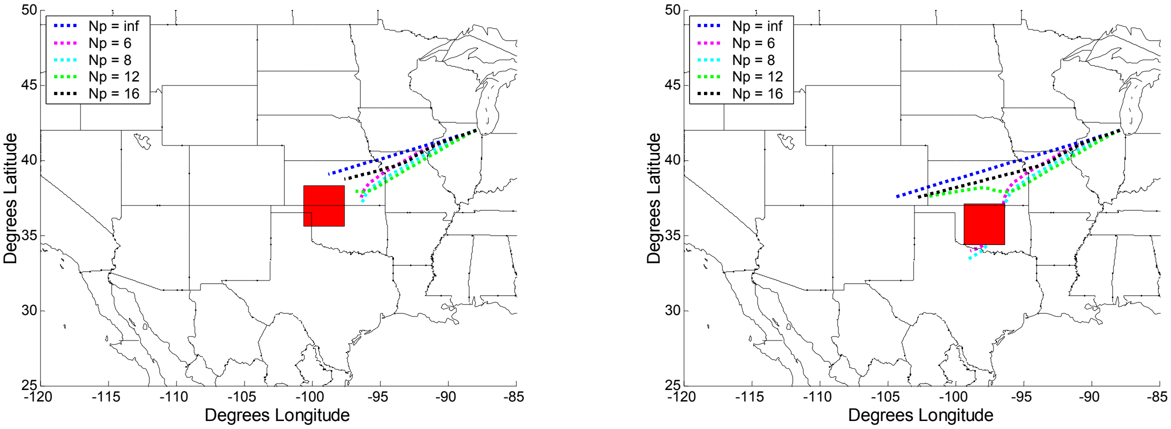

t = 72 min

t = 108 min

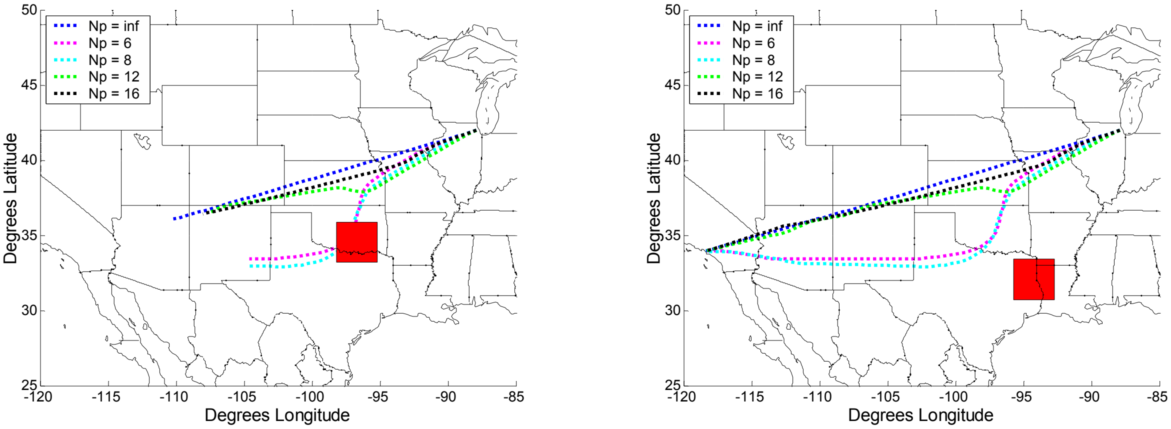

t = 144 min

t = end

Figure 5.17. Comparison of trajectories (N_e = 2, *dt* = 2.4 min) around a medium sized dynamic obstacle with hard avoidance constraints.

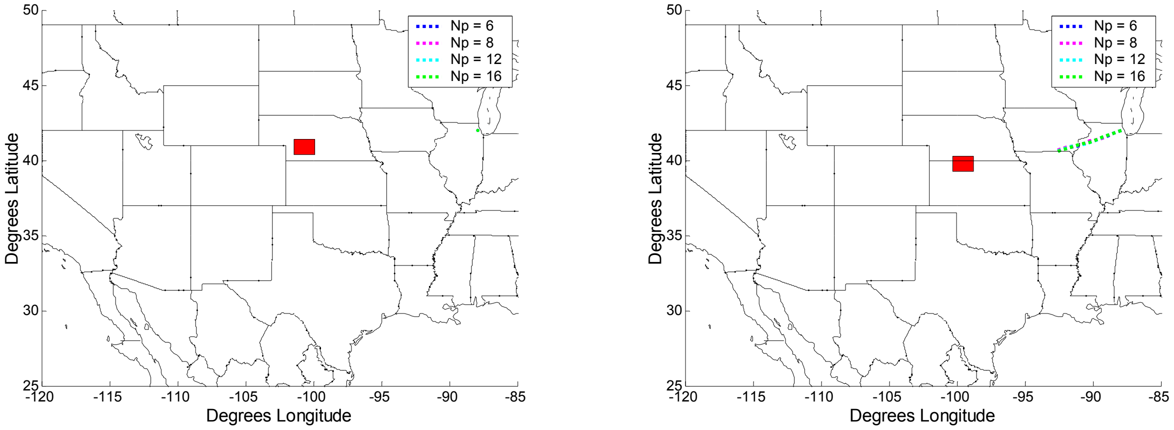

t = 0

t = 36 min

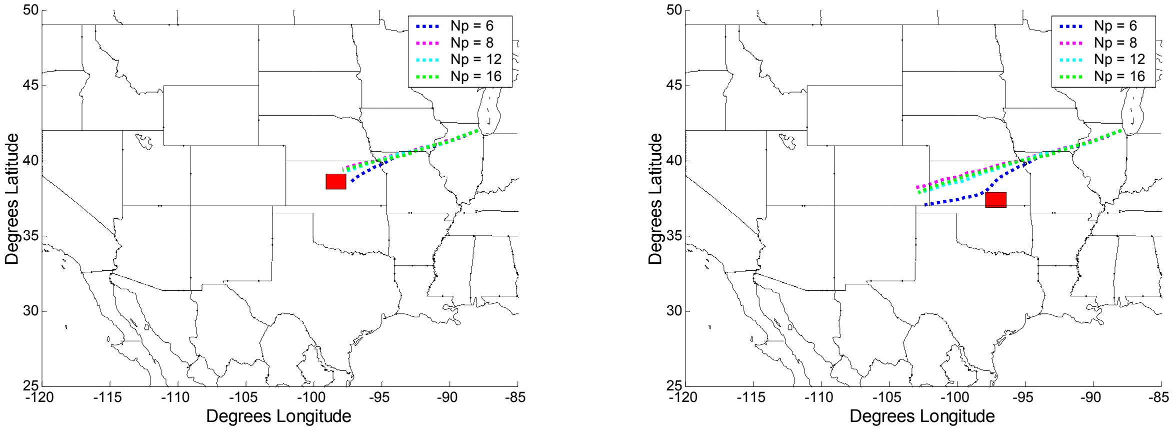

t = 72 min

t = 108 min

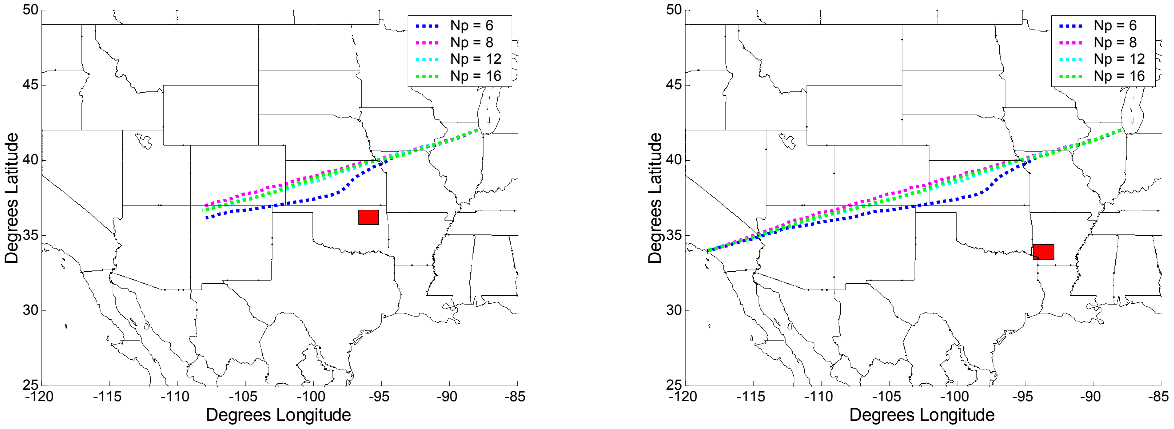

t = 144 min

t = end

Figure 5.18. Comparison of trajectories (N_e = 2, *dt* = 2.4 min) around a small sized dynamic obstacle with hard avoidance constraints.

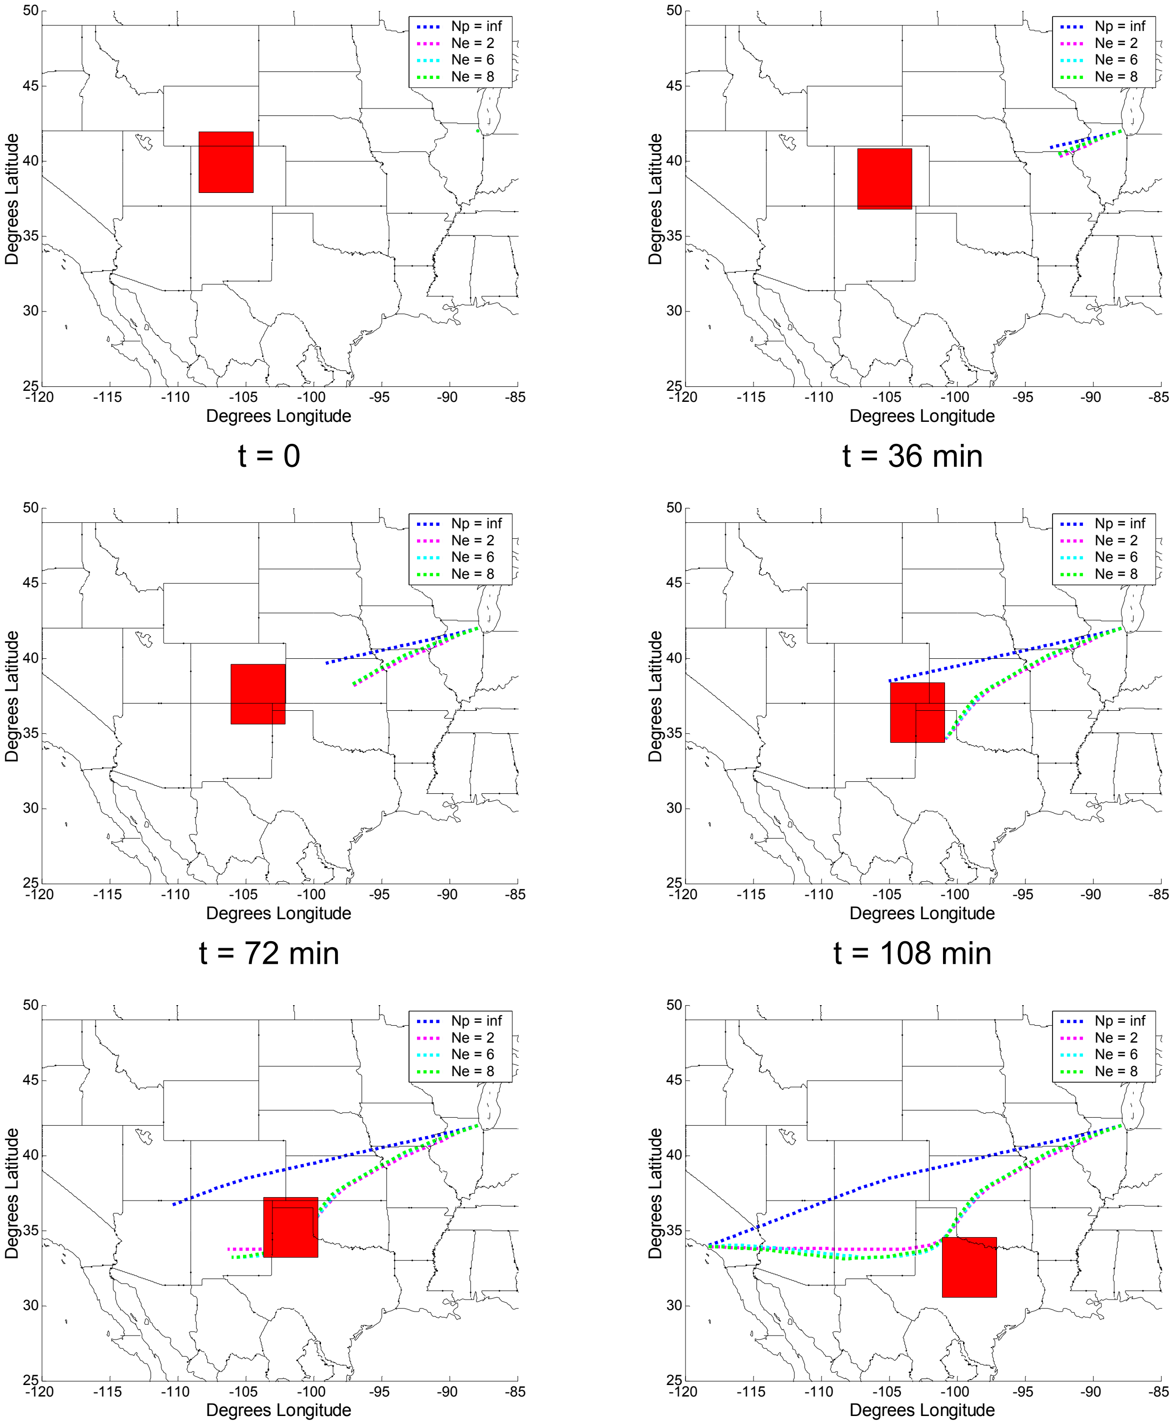

t = 0

t = 36 min

t = 72 min

t = 108 min

t = 144 min

t = end

Figure 5.19. Comparison of trajectories (N_p = 12, dt = 2.4 min) around a large sized dynamic obstacle with hard avoidance constraints.

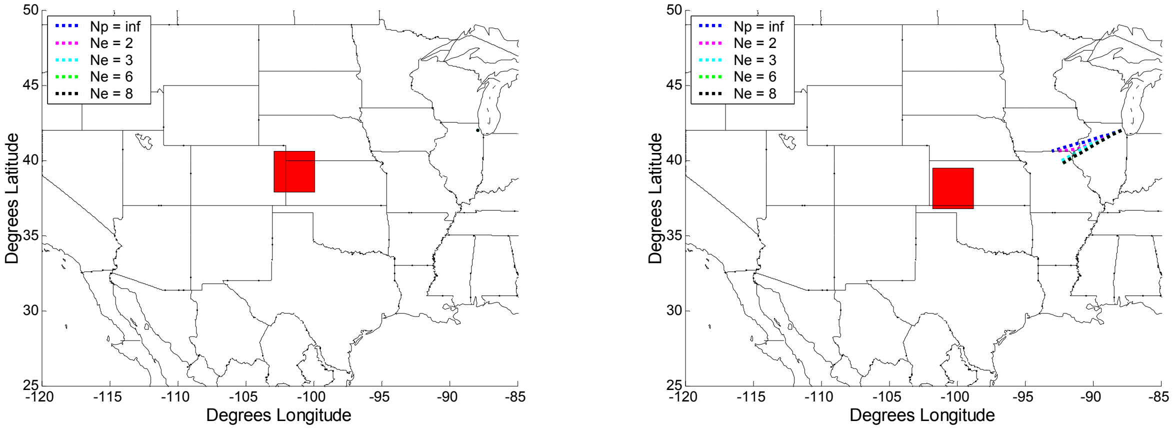

t = 0

t = 36 min

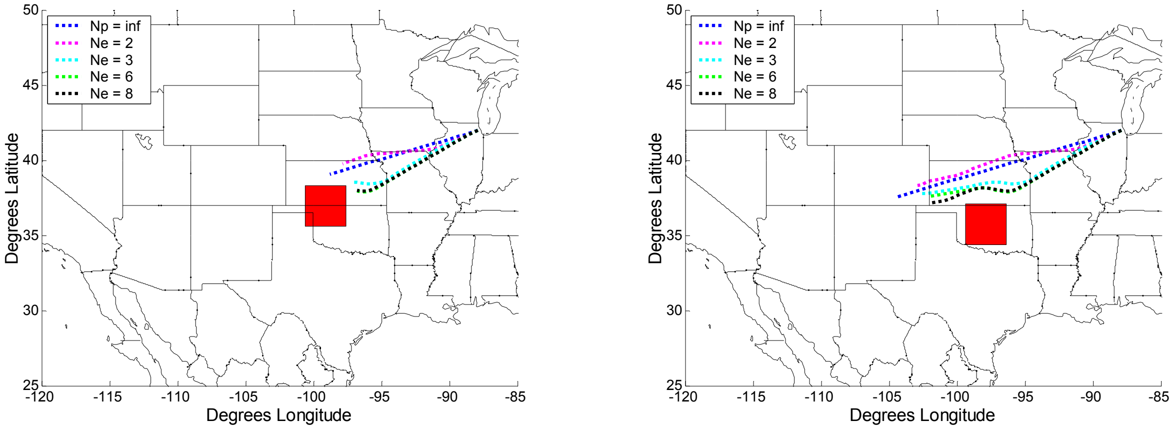

t = 72 min

t = 108 min

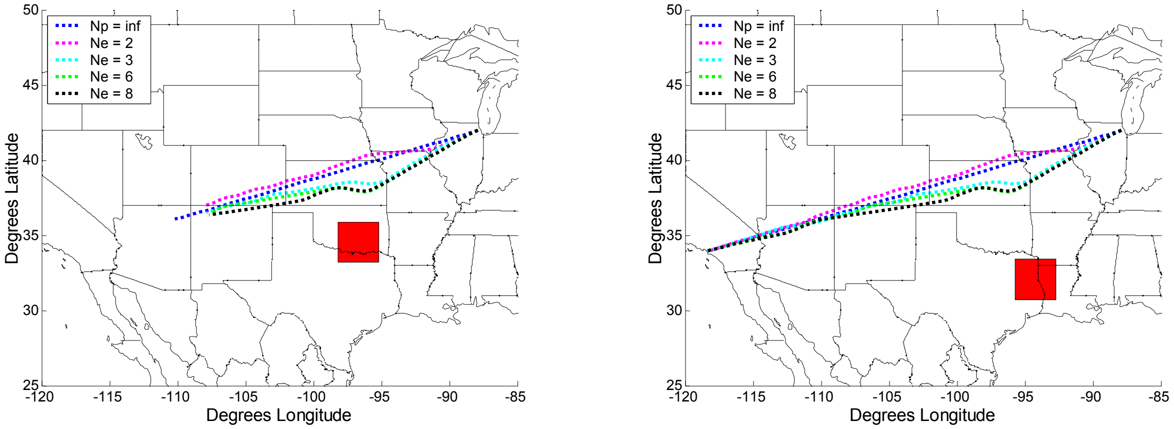

t = 144 min

t = end

Figure 5.20. Comparison of trajectories (N_p = 12, dt = 2.4 min) around a medium sized dynamic obstacle with hard avoidance constraints.

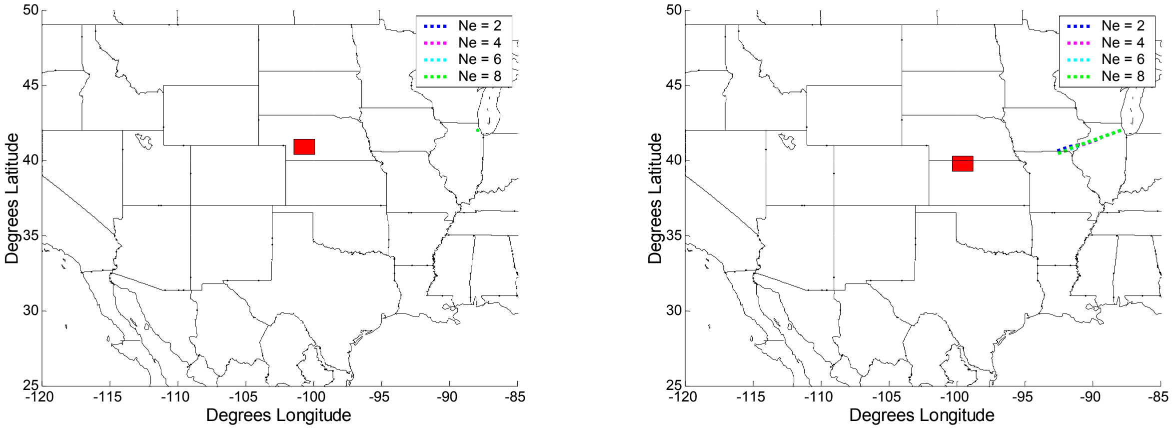

t = 0 t = 36 min

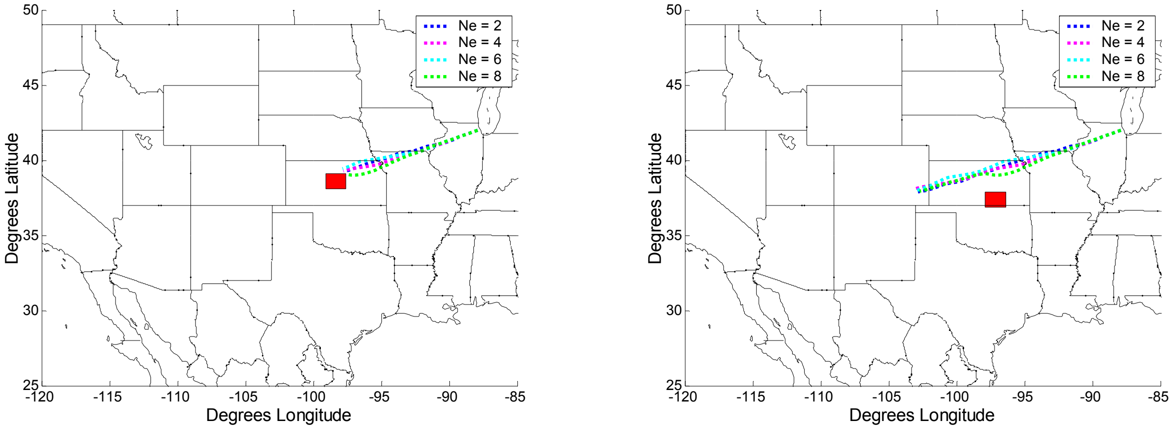

t = 72 min t = 108 min

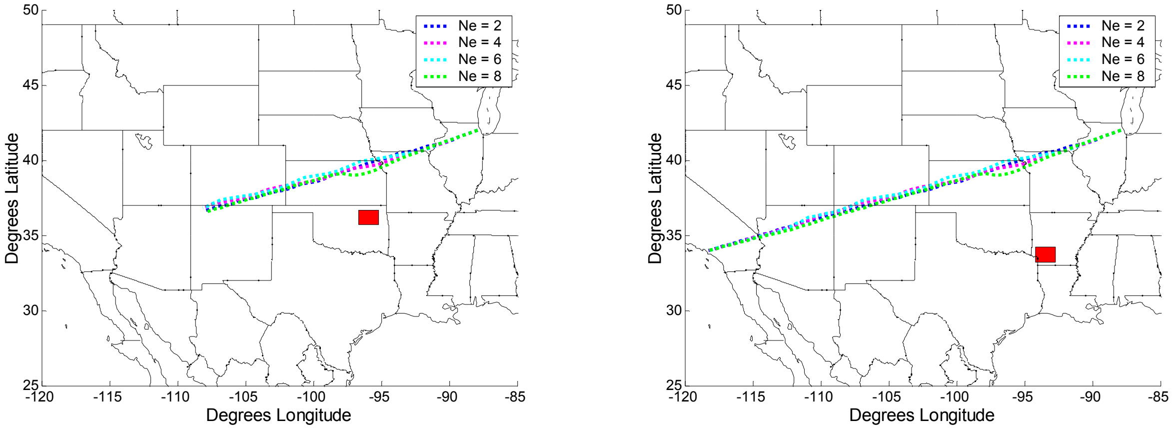

t = 144 min t = end

Figure 5.21. Comparison of trajectories (N_p = 12, dt = 2.4 min) around a small sized dynamic obstacle with hard avoidance constraints.

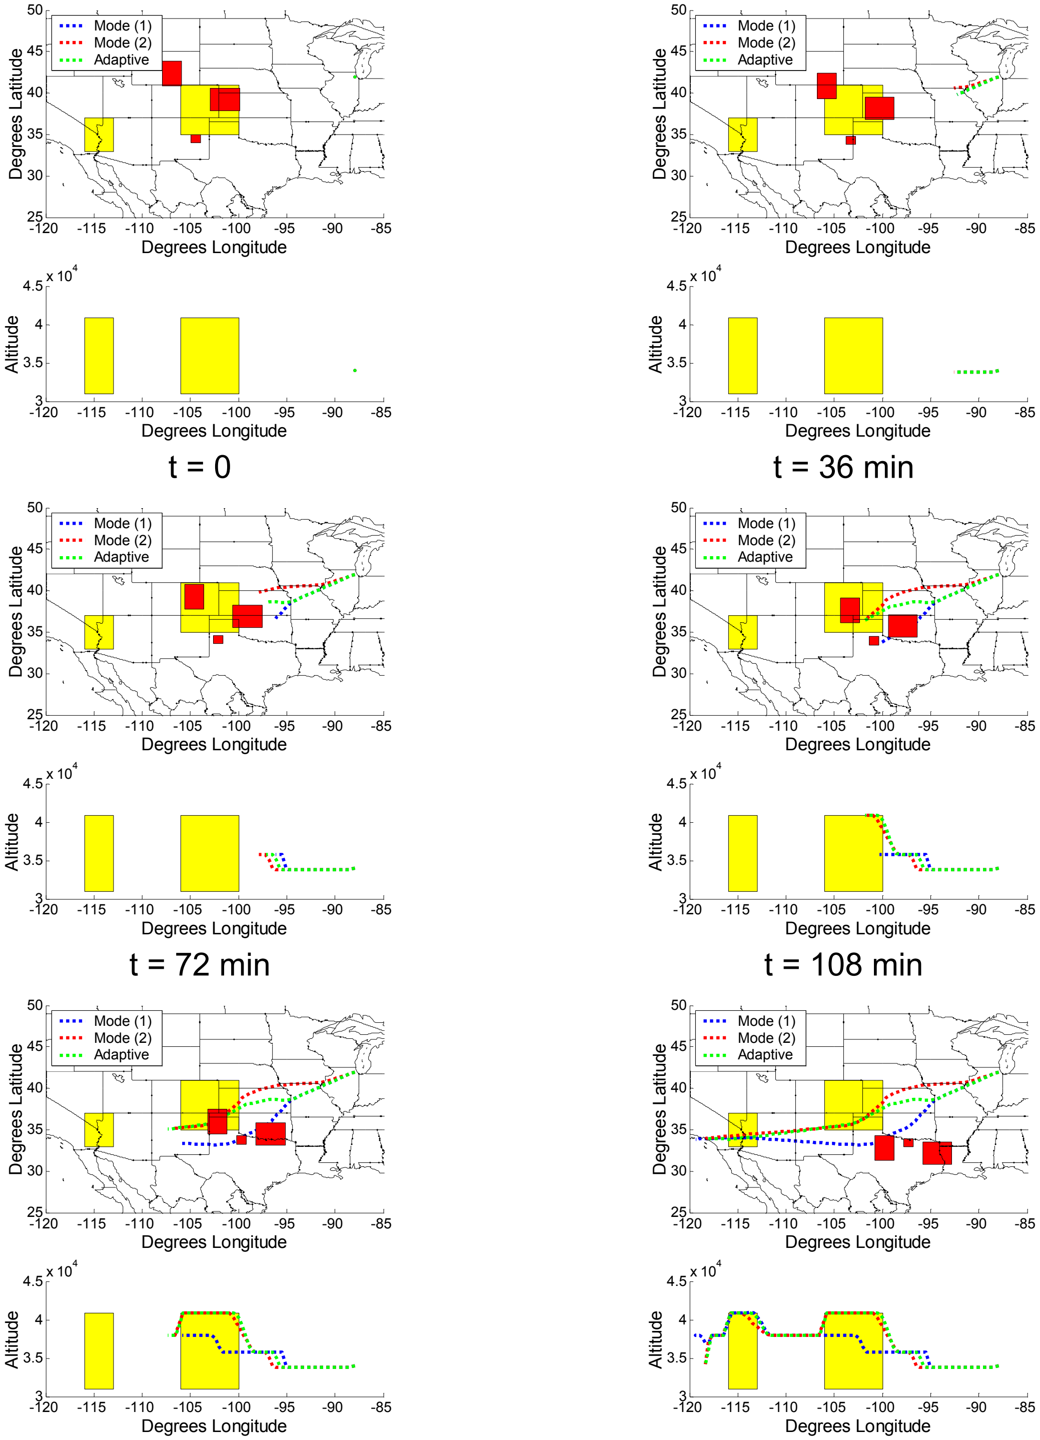

t = 0

t = 36 min

t = 72 min

t = 108 min

t = 144 min

t = end

Figure 5.22. Comparison of the adaptive receding horizon strategy and receding horizon strategies of fixed parameters .

Chapter 6

Practical Results and Additional Applications

6.1 Introduction

This chapter presents applications of the receding horizon path planning algorithm developed in this dissertation. One of the objectives of this research is to develop a method to reduce the environmental impact of aviation. Persistent contrail formation is one of the larger effectors on the environment, and its mitigation is studied as an application of the receding horizon mixed-integer linear programming (RH-MILP) path planner developed in this research. Persistent contrail mitigating trajectories are generated using RH-MILP and studied in depth for a single route example, and then the aggregate effect of the strategy is analyzed using data from multiple days. The results of the persistent contrail mitigation application are compared to existing contrail mitigation techniques in the literature.

Additionally, convective weather is a disturbance that causes a loss of efficiency in the National Airspace System (NAS). Fuel optimal routes around areas of convective weather are generated, and the performance of these trajectories is discussed. The convective weather and persistent contrail scenarios are combined to form a multi-scale path planning problem to generate minimum fuel trajectories for reduced environmental impact. Additional applications of the path planner are introduced, but not quantified.

6.2 RH-MILP for Persistent Contrail Mitigation

The formation of persistent contrails is theorized to be one of the larger effectors on the environment from aviation. Contrails form when the exhaust and entrained air pass through a thermodynamic state that is saturated with respect to water. *Persistent* contrails form when an aircraft creates a contrail in a specific area of the atmosphere called an ice super saturated region (ISSR), where the relative humidity with respect to ice (RHi) is greater than or equal to 100%. In theory, if an aircraft avoids flying in ISSRs, it will not produce persistent contrails. The examples of this section use the atmospheric model discussed in Chapter 2, where the primary source of data is the Rapid Update Cycle (RUC). A single flight example is presented first, followed by the aggregate performance of the path planner, which accounts for data from multiple days. Lastly, the results are compared against existing contrail mitigation strategies in the literature.

6.2.1 Example Persistent Contrail Mitigation Scenario

This example considers a single flight from O'Hare International Airport (ORD) to Los Angeles International Airport (LAX) using atmospheric data from November 17, 2001. The objective of this example was to find a fuel optimal trajectory for this route while minimizing path length in areas containing RHi > 100%. The fuel burn cost was derived from the model presented in Chapter 3.2, and the formulations of Chapter 3.4 describe the dynamical and aircraft performance constraints. Table 6.1 lists the receding horizon parameters and the aircraft performance limitations used in this example.

Figures 6.1 and 6.2 show trajectories overlaid on two-dimensional contour plots of the RHi field at different altitudes. The blue, black, and red trajectories correspond to optimizations with 100%, 50%, and 0% contrail mitigation, respectively. The trajectories were initiated at approximately 34,000 ft, it should be noted that the red and black trajectories are overlaid on each other in Figs. 6.1 and 6.2. The 0% and 50% contrail mitigation trajectories follow an almost identical straight line trajectory from ORD to LAX with the only difference being in the altitude. The 100% contrail mitigation trajectory was forced to adjust its horizontal flight path in addition to its altitude to avoid flight into an area of RHi > 100%.

Figures 6.3 and 6.4 shows the aircraft and contrail mitigation performance associated with the trajectories presented in Figs. 6.1 and 6.2. Figure 6.3(a) shows the velocity profiles of the trajectories, and it is easy to see that the Mach numbers of the 0% and 50% trajectories (red and black respectively) remain constant at roughly 0.78. On the other hand, due to a longer flight path, the 100% trajectory (shown in blue) increases its Mach number in an attempt to arrive in LAX at the same time as the 0% trajectory. Figure 6.3(b) shows the altitude profiles of the trajectories. The 0% trajectory does not change its altitude to avoid contrail formation. The altitude increases on the 0% trajectory because the nominal flight condition changes for the step climb procedure. The 50% and 100% trajectories have the same nominal flight condition as the 0% but adjust their altitude to fly either above or under areas of RHi > 100%. The persistent contrail length produced by each trajectory was normalized by the straight line distance between ORD and LAX. This normalized contrail length was 0.42 for the 0% trajectory, 0.22 for the 50% trajectory, and 0.0 for the 100% trajectory.

Table 6.2 compares the performance of the three trajectories presented by this example. The trajectory with 100% contrail penalty avoided producing any persistent contrails, but it consumed significantly more fuel than the other two trajectories. The 50% contrail penalty trajectory mitigated contrail formation by almost 50% by altering its altitude and only had a slight increase in fuel burn.

6.2.2 Aggregate Route Results

In Chapter 2, it was mentioned that the likelihood of existence of an ISSR depends on the time of day and also the time of year. Figure 6.5 shows the dependence of persistent contrail frequency to the time of year [55, 80]. To understand the complete tradeoff between persistent contrail mitigation and increased fuel burn, this dependency needs to be considered. Therefore, a selection of data from different times of day and times of year were compiled and tested using the path planning algorithm developed in this research. The example uses the same ORD-LAX departure/arrival pair as before. The results of this study are given in Table 6.3. There are 20 days of data, from 4 distinct months at different times of the year. The table presents the path length of persistent contrails for a trajectory with no contrail mitigation logic, the path length of persistent contrails produced with a contrail mitigation penalty in the algorithm, and the increased fuel burn for the persistent contrail mitigation trajectories.

Ideally, a year's worth of data would be collected and analyzed to find the yearly average increase in fuel burn and percent of persistent contrails mitigated, and this is recommended as a future extension of this dissertation. In lieu of 365 days of data, 20 days of data were selected to represent 4 weeks of data at different times of the year in an attempt to capture the seasonal variation in contrail production. It was found that the strategy presented in this dissertation was able to mitigate 58% of persistent contrails with a 0.48% increase in fuel burn. Of the 20 days of data, 9 were found to be conducive to persistent contrail formation, which is in the ballpark of the data presented in Fig. 6.5.

6.2.3 Assessment of Operational Strategies

Operational strategies for persistent contrail mitigation have been studied in the past. Klima showed that non-optimal contrail mitigation strategies can reduce persistent contrail coverage [21]. For the case of individual rerouting (each aircraft is rerouted independently), persistent contrails were reduced 65%-80%. For the case of weekly rerouting (custom routes are changed on a weekly basis), persistent contrails were reduced 55%-85% with a 1%-2% increase in operational costs. Other strategies proposed by Klima include routing aircraft away from the humid tropopause, flying more fuel efficient routes, and choosing a more northerly route for transatlantic flights [21].

Another strategy takes advantage of the geometry of the regions of ice super saturation. Because regions of ice supersaturation are very thin, it has been proposed that altitude change is the best method to avoid contrail production. Mannstein et al. showed that small changes in aircraft altitude can significantly reduce the impact of contrails [22]. Williams and Noland also assessed the viability of altitude changes on contrail formation [23]. Fichter et al. found that contrail coverage could be reduced approximately 45% by flying 6000 feet lower on average with a 6% penalty in fuel burn [24].

This research is the first strategy to use optimization methods to determine the best trajectory for contrail mitigation and fuel conservation. However, the drawback to this strategy is that it requires a large infrastructure and system improvement to the NAS for implementation, whereas some of the strategies in the literature do not. Therefore, the optimal trajectories of this research can help determine the most effective method that can be implemented easily. A finding of this research is that the optimal method for mitigating and/or avoiding persistent contrail formation is a simple attitude adjustment. This corroborates a subset of the existing literature [22-24] on operational methods for persistent contrail mitigation. Some of these methods recommend placing altitude restrictions, or artificial ceilings, on flights that would create persistent contrails [23, 24]. However, the infrequency of ice super saturated regions lends to the idea that altitude restrictions are not the most efficient solution. For example, the study of Fichter et al. found that persistent contrails can be reduced 45% by flying 6000 feet lower on average with a 6% penalty in fuel burn [24], which is a much larger performance penalty than showed by this research. The study that most closely aligns with the results of this research is the work of Mannstein et al, which calls for real-time altitude adjustments around ice super saturated regions once the region was sensed by the aircraft [22].

6.3 RH-MILP for Convective Weather Avoidance

Thunderstorms are a leading cause of delay in the NAS, where the delays are a result of ground holds, in-flight holds, and general rerouting of aircraft around storms. The increased flight time corresponding to both in-flight holds and rerouting leads to increased environmental impact due to extra fuel burn emissions. This section presents an application of the path planning algorithm developed in this research to thunderstorm avoidance. The path planner generates minimum fuel routes around areas of convection that would hopefully alleviate some of the strain on the NAS due to thunderstorms.

Pilots have been observed to change their route to avoid thunderstorms with a VIP rating of 3 or greater. In this research, areas with a VIP ≥ 3 were represented as cuboids, and flight through these regions was precluded with hard avoidance constraints. The dynamics of a thunderstorm are very complicated and can be extremely hard to predict. The path planner used the procedure outlined in Chapter 2 to account for the thunderstorm dynamics, where the general movement of the thunderstorm was predicted based on an extrapolation of the previous two time steps, and a margin of safety was implemented depending on the level of unpredictability of the storm. This section presents an example of a single flight which encounters thunderstorms en-route. The performance of the path planning algorithm is analyzed for the results of this application. Lastly, the performance of the path planner is observed in a scenario that includes both persistent contrail mitigation and convective weather avoidance.

6.3.1 Example Convective Weather Avoidance Scenario

This example considers a single flight from O'Hare International Airport (ORD) to Los Angeles International Airport (LAX) using atmospheric data from May 7, 2008. The objective of this example was to find a fuel optimal trajectory for this route while avoiding flight through a storm with a VIP level three or greater. The fuel burn cost was derived from the model presented in Chapter 3.2, and the formulations of Chapter 3.4 describe the dynamical and aircraft performance constraints. The receding horizon parameters are the same as Table 6.1 except that the execution horizon was set to 2 waypoints.

Figure 6.6 shows the shape of the flight path in the presence of en-route convective weather, where each diagram in Fig. 6.6 corresponds to a different time. The thunderstorms are represented in Fig. 6.6 as green and yellow shaded areas, where a green area corresponds to VIP < 3 and a yellow area depicts an area of VIP ≥ 3. There is a map of the United States in the background to provide perspective, and the trajectory is shown by a red dashed line. At the time of departure, the thunderstorms are located in a line extending from western Kansas to central Wisconsin, and the general movement of the storms is to the east. It is apparent in Fig. 6.6(b) that the trajectory is going to encounter the thunderstorms, and adjusts its route to pass to the south of the storms, as is shown in Fig. 6.6(c). As is seen in Fig. 6.6(f), the path planner successfully negotiated the disturbance of convective weather.

6.3.2 Combined Strategy for Reduced Environmental Impact

This example considers a single flight from O’Hare International Airport (ORD) to Los Angeles International Airport (LAX) using relative humidity data from November 17, 2001, and convective weather data from May 7, 2008. The objective of this example is to find a fuel optimal trajectory for this route while avoiding flight through a storm with a

VIP level three or greater, and assigning a penalty to flight through an area of RHi ≥ 100%. The fuel burn cost is derived from the model presented in Chapter 3.2, and the formulations of Chapter 3.4 describe the dynamical and aircraft performance constraints. This trajectory is generated using the adaptive strategy described in Chapter 5.

Figures 6.7-6.12 show the shape of the flight path in the presence of en-route convective weather and persistent contrail mitigation constraints, where each figure corresponds to a different time. The thunderstorms are represented as green and yellow shaded areas, where a green area corresponds to VIL < 3 and a yellow area depicts an area of VIL ≥ 3. Areas of ice super saturation are shown by red shading, based on the RHi field at approximately 34,000 feet. Three trajectories of different persistent contrail formation penalty are compared, and the results are shown in Table 6.4. All three trajectories follow the same path in the horizontal plane, which is the same as the horizontal path generated in Fig. 6.6. The differences between the trajectories can be seen in the vertical plane, where the different persistent contrail weightings cause the path planner to adjust altitude to control the amount of path length within an area of ice super saturation. The trajectory with no penalty on contrail formation follows the nominal step climb profile, whereas the trajectory with the persistent contrail penalty tuned to 50% avoidance continues its climb at the first step climb point to a higher altitude above the area of RHi ≥ 100%. The 100% persistent contrail avoidance trajectory starts its climb over the ice super saturation regions much earlier, eliminating the majority of the persistent contrails. The 0% avoidance trajectory was used as a baseline to examine the tradeoff between persistent contrail mitigation and increased fuel burn. The 0% avoidance trajectory theoretically created persistent contrails with a combined length of 763 nm, compared to 544 nm for the 50% avoidance trajectory and 80 nm for the 100% avoidance trajectory. The 50% avoidance trajectory burned 0.9% more fuel than the 0% trajectory, and the 100% avoidance trajectory burned 2.1% more fuel than the baseline.

6.4 Additional Applications

In addition to the persistent contrail mitigation and persistent contrail avoidance applications presented in this chapter, there are other potentially useful applications of this research. The three additional applications introduced in this section are sonic boom mitigation, turbulence avoidance, and aircraft icing avoidance. The sonic boom is a design driver for the development of supersonic transport because of possible increases in noise pollution along the flight path, especially over land routes. Turbulence and aircraft icing have recently contributed to accidents in air transportation. This section outlines how the path planning algorithm developed in the research could be applied to these issues in aviation.

6.4.1 Sonic Boom Impact Mitigation

Traditionally, noise impacts are concentrated to areas very close to airports, where the sound of arriving and departing aircraft affect the nearby population. However, with a new emphasis on supersonic flight, the noise generated from a sonic boom will make noise pollution a more widespread impact [81]. One strategy proposed to mitigate the impact of the sonic boom is to avoid repeated overflight of the same region, which essentially attempts to limit an area's exposure to sonic booms [82]. For example, consider an arrival-departure pair of Teterboro, NJ, and San Jose, CA. If the optimal roundtrip is flown 4 times per day, then the same points along the overland trajectory will be subjected to the sustained sonic boom impact 8 times per day. This number can be reduced by restricting flight over the previously flown trajectories of the day.

Sonic boom mitigation can be accomplished using the path planning algorithm described in this research by coding the "boom carpet" region of a trajectory with avoidance constraints, and assigning a penalty to subsequent flights for overlapping boom carpets. A boom carpet is the area of ground where the sonic boom is audible. Figure 6.13 shows a simulation of this scenario, where the first trajectory tracks a great-circle trajectory from Teterboro to San Jose, and the return flight minimizes the overlap of the boom carpet regions, which diminishes boom carpet areas based on the previous overflight.

6.4.2 Turbulence Avoidance

Turbulence is a common occurrence on commercial flights, and is widely recognized as the primary cause of injuries to the passengers and crew onboard commercial aircraft. There are many causes of turbulence, such as thunderstorms, mountain waves, and clear air turbulence. In addition, aircraft may encounter turbulence anywhere, including in remote areas where the turbulence forecasts lack the necessary resolution to prevent an incident due to turbulence. The recent crash of Air France flight 447 on June 1, 2009 has highlighted the importance of avoiding en-route turbulence. Although the cause of the crash is still uncertain, it is believed that the aircraft encountered severe turbulence and crashed off the cost of Brazil. The crash occurred in a remote area of the ocean, where there is limited weather information available to pilots.

The National Center for Atmospheric Research (NCAR) has started research to develop a global forecast system to predict areas of turbulence [83]. Figure 6.14 shows some results of recent NCAR research [84], specifically a prediction of the turbulence encountered by Air France 441 on June 1, 2009. NCAR is developing the forecast for implementation into the Next Generation Air Transportation System's planned 4-D aviation weather database, and these data are a natural application of the path planning algorithm developed in this dissertation. The areas of predicted turbulence can be coded with hard avoidance constraints to preclude any feasible trajectory from the area, and provide an additional margin of safety to commercial flights.

6.4.3 Icing Avoidance

Aircraft icing is widely recognized as a significant hazard to aircraft operations. During the years 1990 to 2005, 35,317 aircraft accidents and incidents were reported in the National Transportation Safety Board (NTSB) accident database [85]. Out of these 35,317 accidents and incidents, 803 were related to flight into icing conditions. Of the accidents involving structural icing, 18 occurred during part 121 (commercial) operations, and 91 occurred during part 135 (on-demand) operations. According to the database, the effects of icing are not limited to smaller aircraft. Transport category aircraft including mid-size jet aircraft, regional jets, and turboprops have all experienced problems with icing. The most recent fatal accidents involving a transport category aircraft occurred on February 12, 2009 and January 9, 1997. On February 12, 2009, Colgan 3407 crashed on final approach to Buffalo Niagara International Airport due to an inadvertent stall partially caused by flight into icing conditions. On January 9, 1997, Comair 3272 entered icing conditions and departed controlled flight on approach to Detroit Metropolitan Wayne County Airport.

The path planning algorithm developed in this research may be applied to icing avoidance given knowledge of where the icing conditions exist. The Current Icing Product (CIP) combines data from multiple sources to create a detailed three-dimensional hourly forecast of the potential for icing and supercooled large droplet (SLD) conditions [86]. Figure 6.15 shows an output of the CIP that could be used to model the avoidance regions of the path planner. Note that the SLD conditions could be treated with hard constraints, and the other areas of icing could be treated as soft constraints with a penalty function.

6.5 Summary

This chapter presented practical applications and simulation results of the path planning algorithm developed in this dissertation. The algorithm proved to be successful in negotiating an environment with a persistent contrail formation penalty, as well as convective weather avoidance constraints. The performance of the path planning algorithm in the persistent contrail mitigation scenario was analyzed for multiple days to investigate the operational feasibility of persistent contrail mitigation. The trajectories generated with the path planner were compared to existing non-optimal contrail mitigation strategies to seek a best near term solution for persistent contrail mitigation.

Additional applications of the path planning algorithm were introduced and remain to be studied in detail in future work.

Table 6.1. Receding horizon parameters and aircraft performance limits.

Number of steps in planning horizon (N_p)	12
Number of steps in the execution horizon (N_e)	6
Time step size (Δt)	3 min
Maximum en-route velocity (V_{max})	470 knots
Minimum en-route velocity (V_{min})	417 knots
Maximum altitude (z_{max})	42,000 ft
Minimum altitude (z_{min})	28,000 ft

Table 6.2. Comparison of receding horizon trajectory performance.

	Contrail Penalty 0%	Contrail Penalty 50%	Contrail Penalty 100%
Max. Velocity	454.1 knots	454.1 knots	467.3 knots
Avg. Velocity	450.6 knots	450.6 knots	464.8 knots
Total Fuel Burn	20,431 lbs	20,734 lbs	21,695 lbs
Flight Time	3.55 hrs	3.55 hrs	3.60 hrs
Normalized Contrail Length	0.42	0.22	0.00

Table 6.3. Performance comparison of trajectories generated with data from different days.

Date:	11/14/2001	11/15/2001	11/16/2001	11/17/2001	11/18/2001
Path length of persistent contrails (no mitigation) (nm)	0	0	264	375	132
Path length of persistent contrails (with mitigation) (nm)	0	0	87	202	65
Increase fuel burn (%)	0	0	0.6%	1.5%	0.4%
Date:	1/12/2002	1/13/2002	1/14/2002	1/15/2002	1/16/2002
Path length of persistent contrails (no mitigation) (nm)	0	402	176	214	0
Path length of persistent contrails (with mitigation) (nm)	0	230	38	55	0
Increase fuel burn (%)	0	0.8%	1.2%	1.3%	0
Date:	5/12/2002	5/13/2002	5/14/2002	5/15/2002	5/16/2002
Path length of persistent contrails (no mitigation) (nm)	129	0	0	0	64
Path length of persistent contrails (with mitigation) (nm)	55	0	0	0	23
Increase fuel burn (%)	0.5%	0	0	0	0.4%
Date:	9/12/2002	9/13/2002	9/14/2002	9/15/2002	9/16/2002
Path length of persistent contrails (no mitigation) (nm)	0	168	0	0	0
Path length of persistent contrails (with mitigation) (nm)	0	50	0	0	0
Increase fuel burn (%)	0	0.8%	0	0	0

Table 6.4. Comparison of receding horizon trajectory performance.

	Contrail Penalty 0%	Contrail Penalty 50%	Contrail Penalty 100%
Max. Velocity	464.3 knots	464.3 knots	464.3 knots
Avg. Velocity	453.6 knots	453.6 knots	453.6 knots
Total Fuel Burn	21,367 lbs	21,559 lbs	21,816 lbs
Flight Time	3.65 hrs	3.65 hrs	3.66 hrs
Normalized Contrail Length	0.48	0.32	0.04

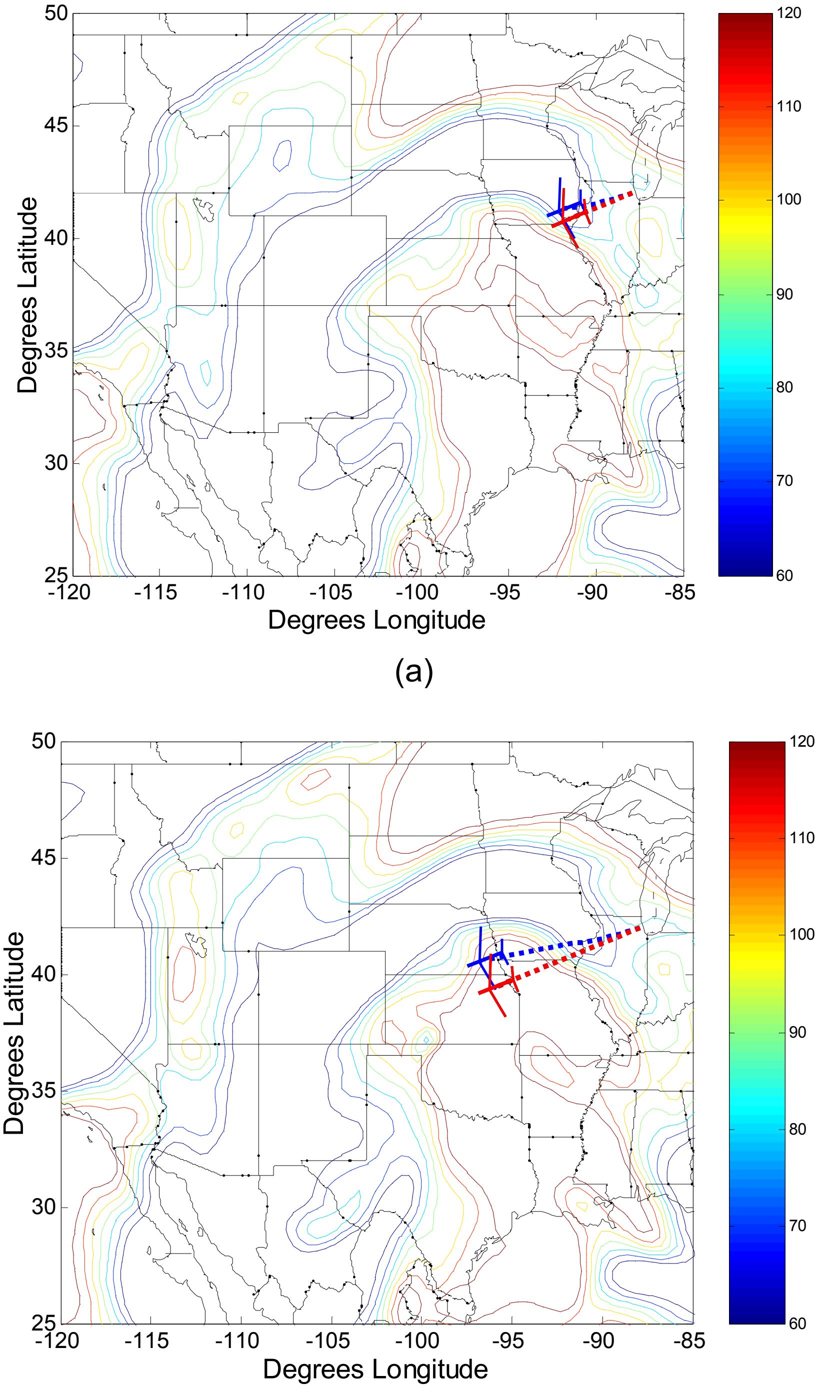

(a)

(b)

Figure 6.1. Fuel optimal trajectories overlaid on contour plots of the Rhi field at times of (a) 1730Z, (b) 1800Z. A map showing the boundaries of North America is in the background. The blue line corresponds to 100% avoidance, the black line to 50% avoidance, and the red line to 0% avoidance.

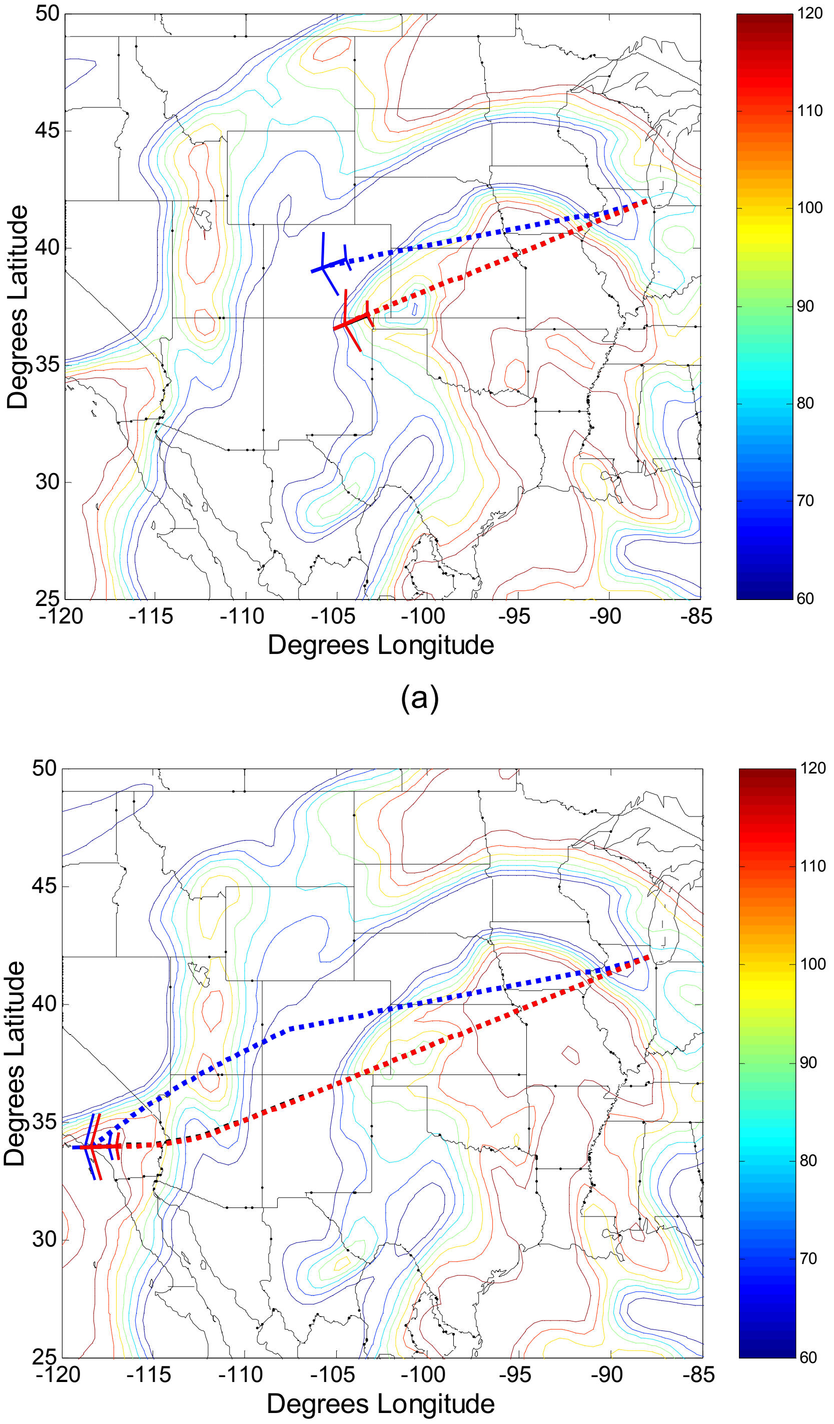

(a)

(b)

Figure 6.2. Fuel optimal trajectories overlaid on contour plots of the Rhi field at times of (a) 1900Z, and (b) 1940Z. A map showing the boundaries of North America is in the background. The blue line corresponds to 100% avoidance, the black line to 50% avoidance, and the red line to 0% avoidance.

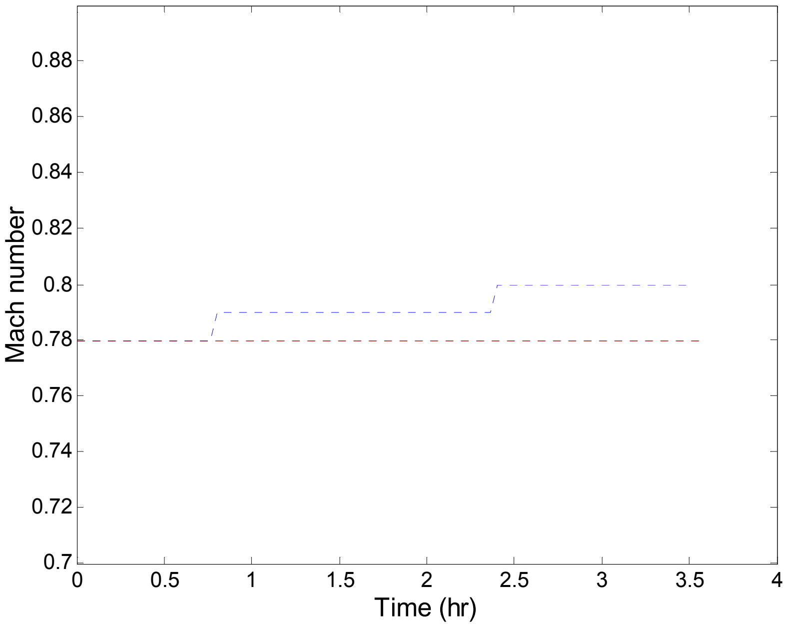

(a)

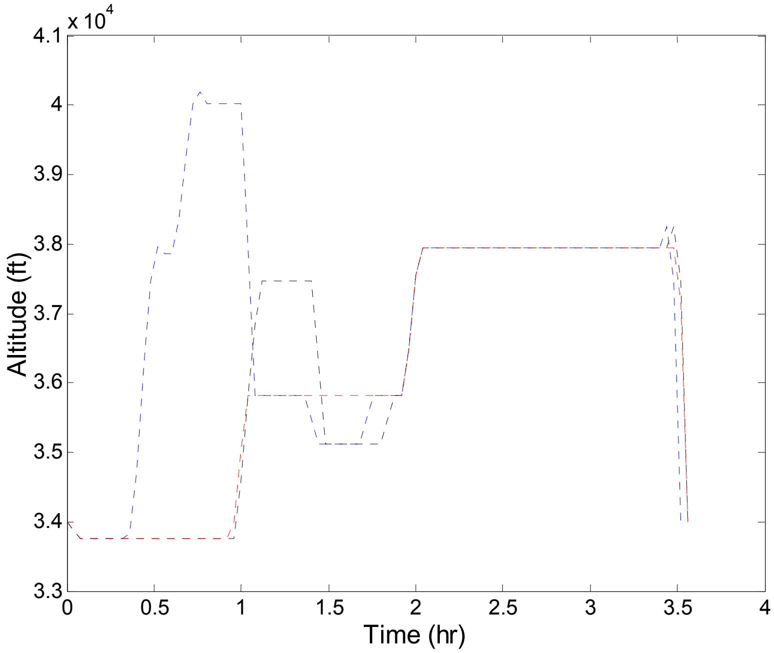

(b)

Figure 6.3. Aircraft performance: (a) velocity time history, (b) altitude time history. The blue line corresponds to 100% avoidance, the black line to 50% avoidance, and the red line to 0% avoidance.

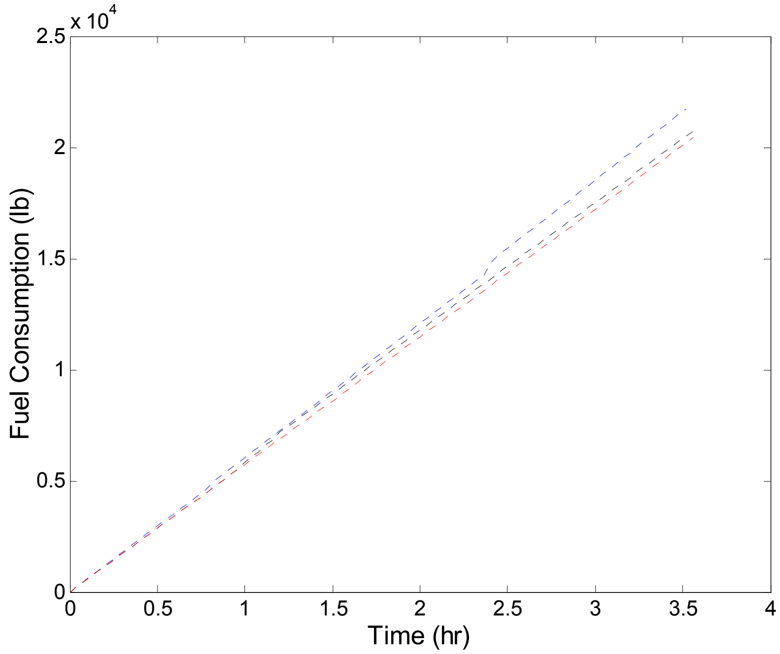

(a)

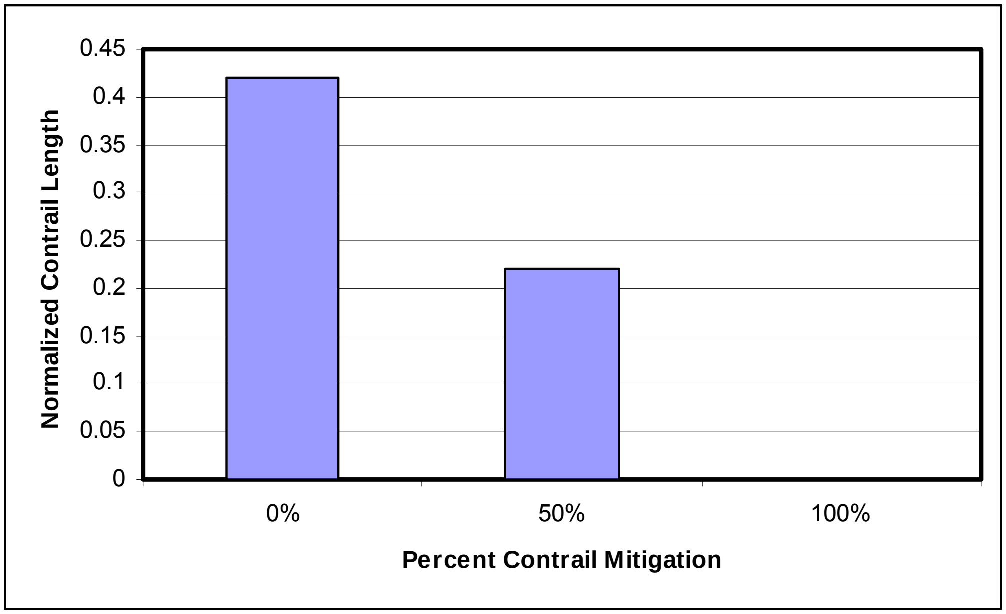

(b)

Figure 6.4. Aircraft performance: (a) fuel burn time history, (b) persistent contrail formation bar graph. The blue line corresponds to 100% avoidance, the black line to 50% avoidance, and the red line to 0% avoidance.

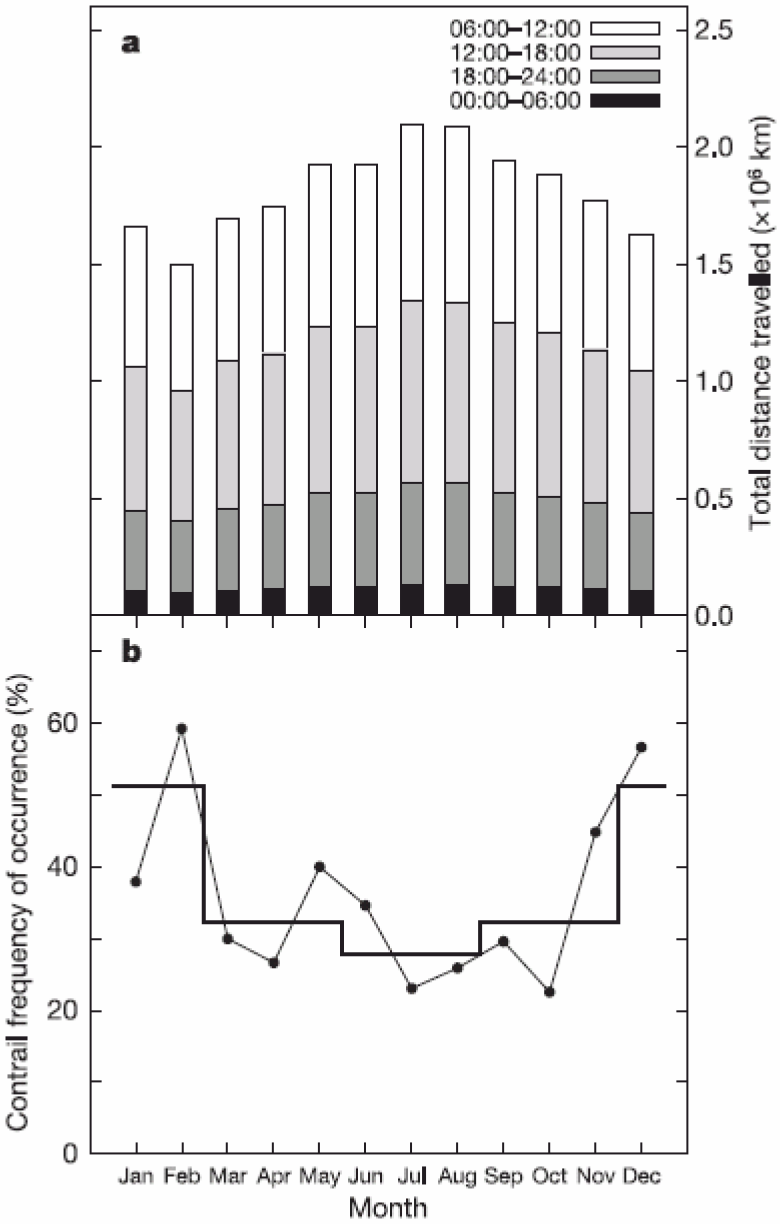

Figure 6.5. Yearly cycle of total distance traveled (a) and contrail frequency (b). The different shading in (a) indicates different times of day, and the heavy black line in (b) shows the seasonal mean [55].

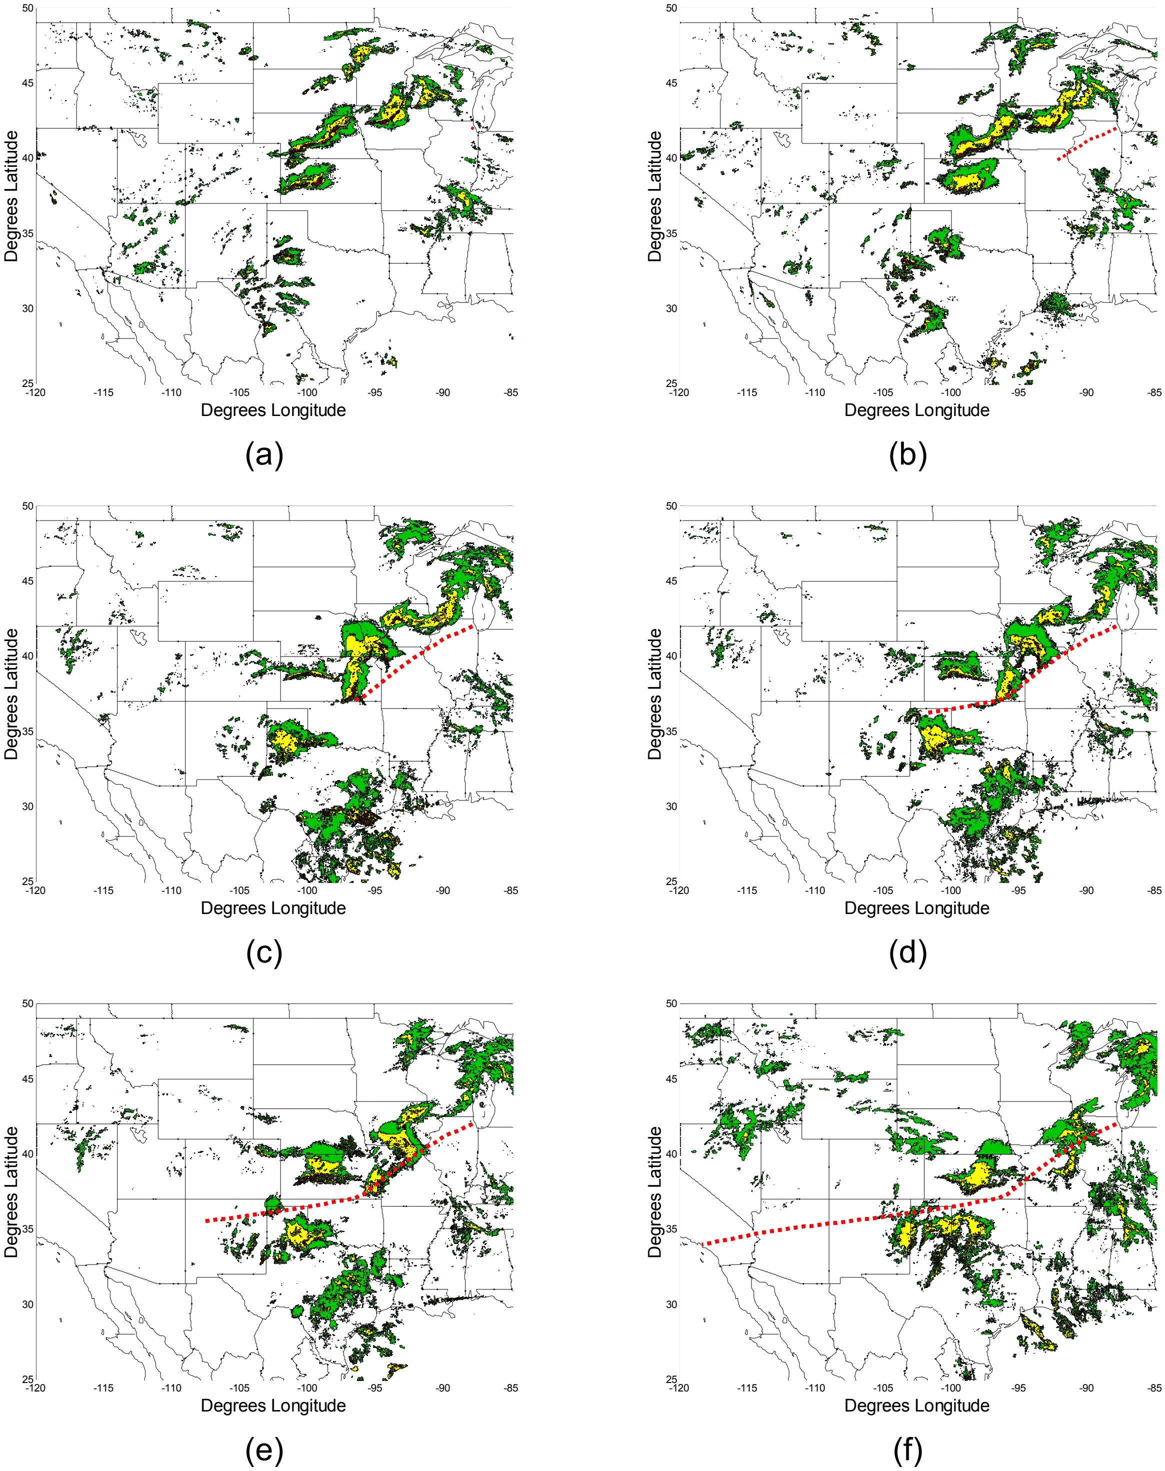

(a) (b)

(c) (d)

(e) (f)

Figure 6.6. Receding horizon trajectory for convective weather avoidance. This trajectory was generated with a planning horizon of 12 and an execution horizon of 2.

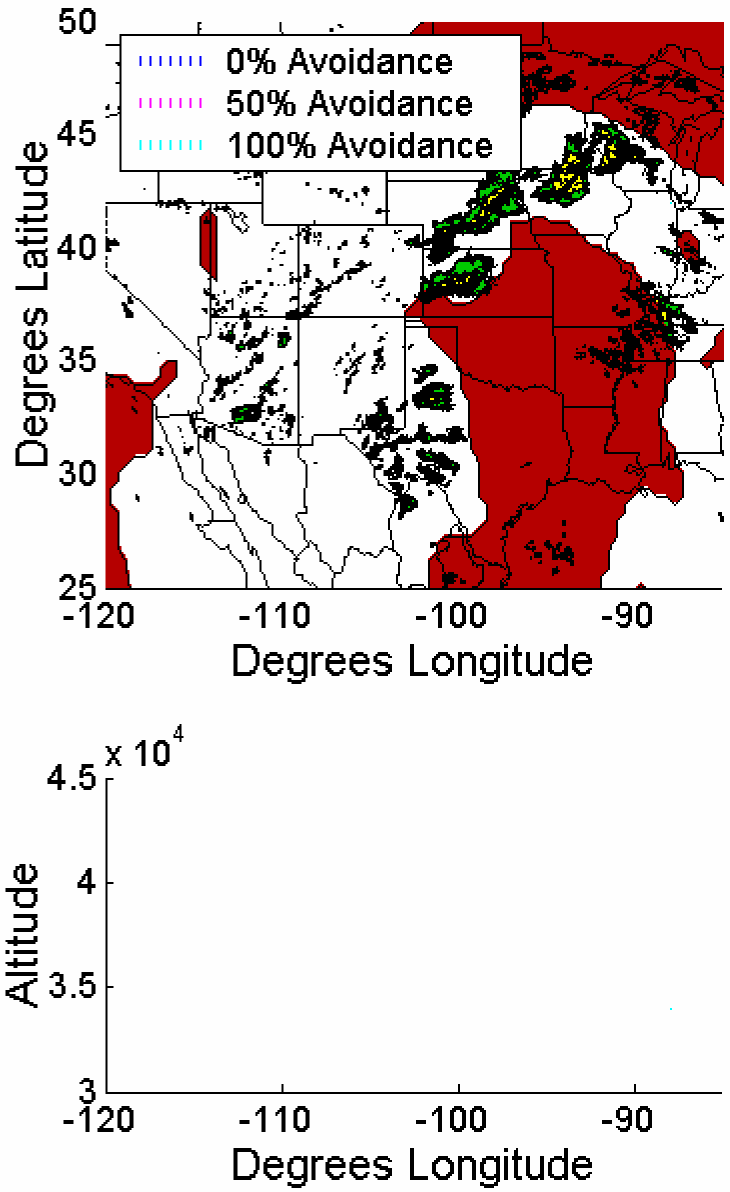

Figure 6.7. Receding horizon trajectories for a scenario involving both persistent contrail mitigation and convective weather avoidance (t = 0 min).

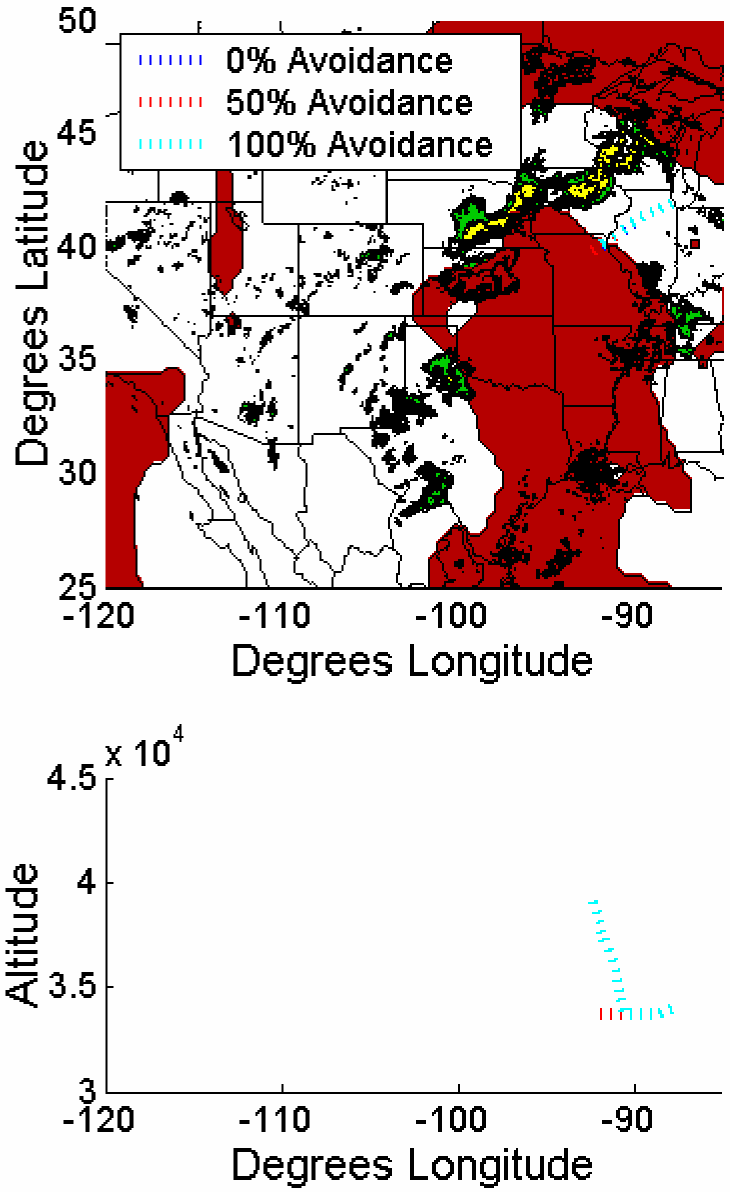

Figure 6.8. Receding horizon trajectories for a scenario involving both persistent contrail mitigation and convective weather avoidance (t = 36 min).

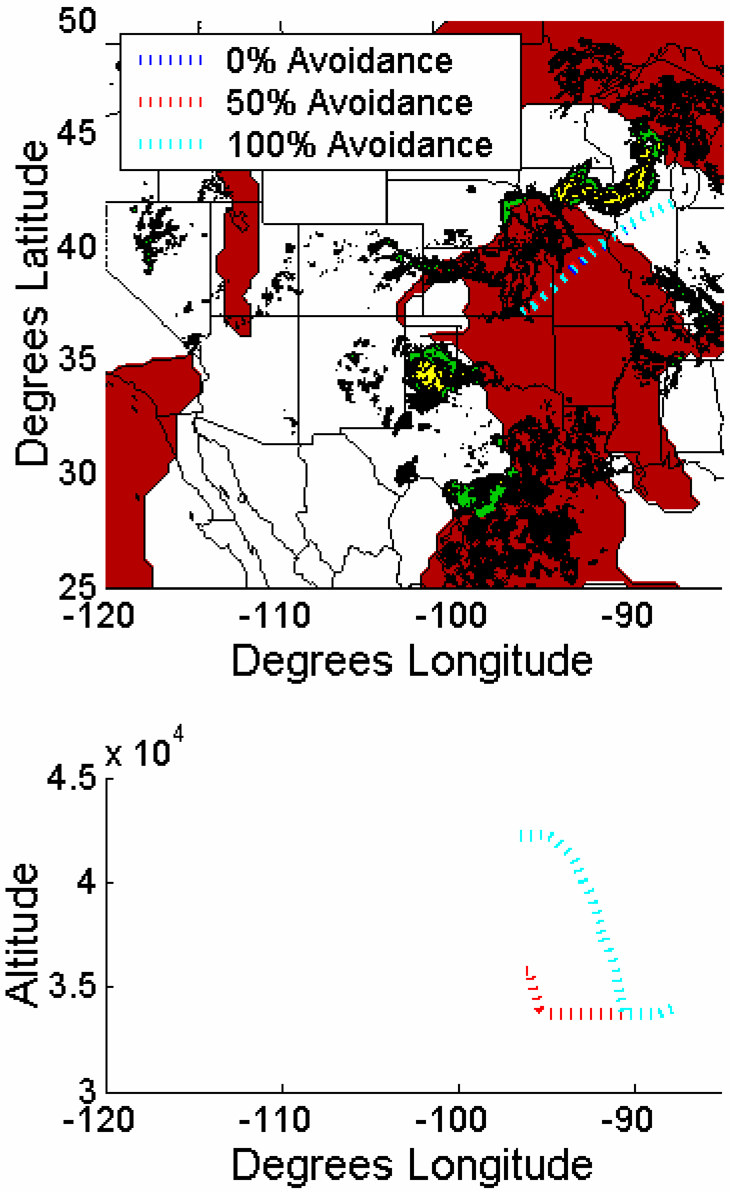

Figure 6.9. Receding horizon trajectories for a scenario involving both persistent contrail mitigation and convective weather avoidance (t = 72 min).

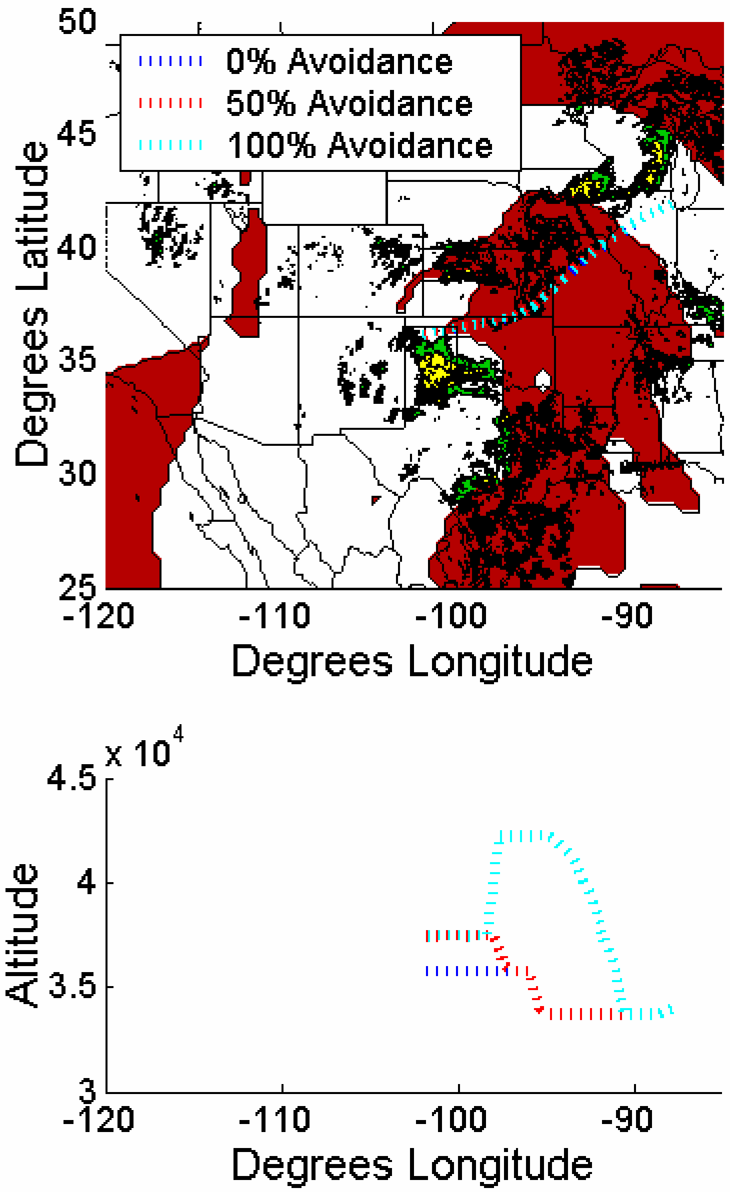

Figure 6.10. Receding horizon trajectories for a scenario involving both persistent contrail mitigation and convective weather avoidance (t = 108 min).

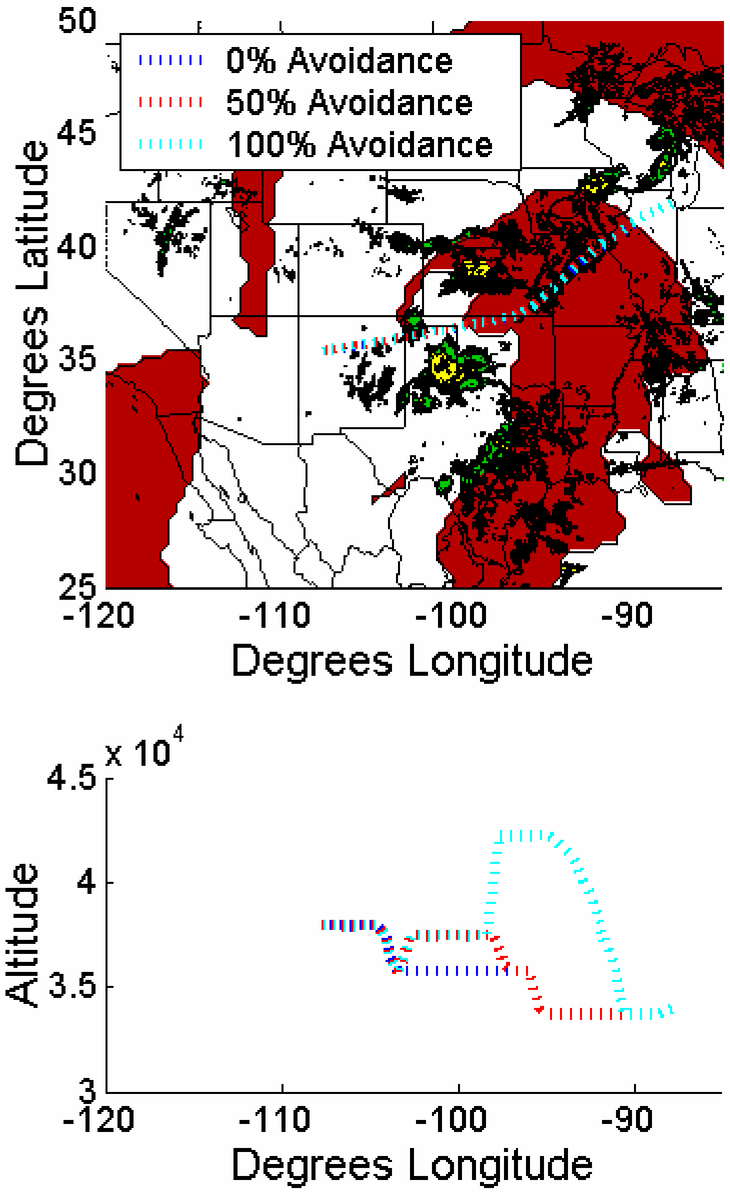

Figure 6.11. Receding horizon trajectories for a scenario involving both persistent contrail mitigation and convective weather avoidance (t = 144 min).

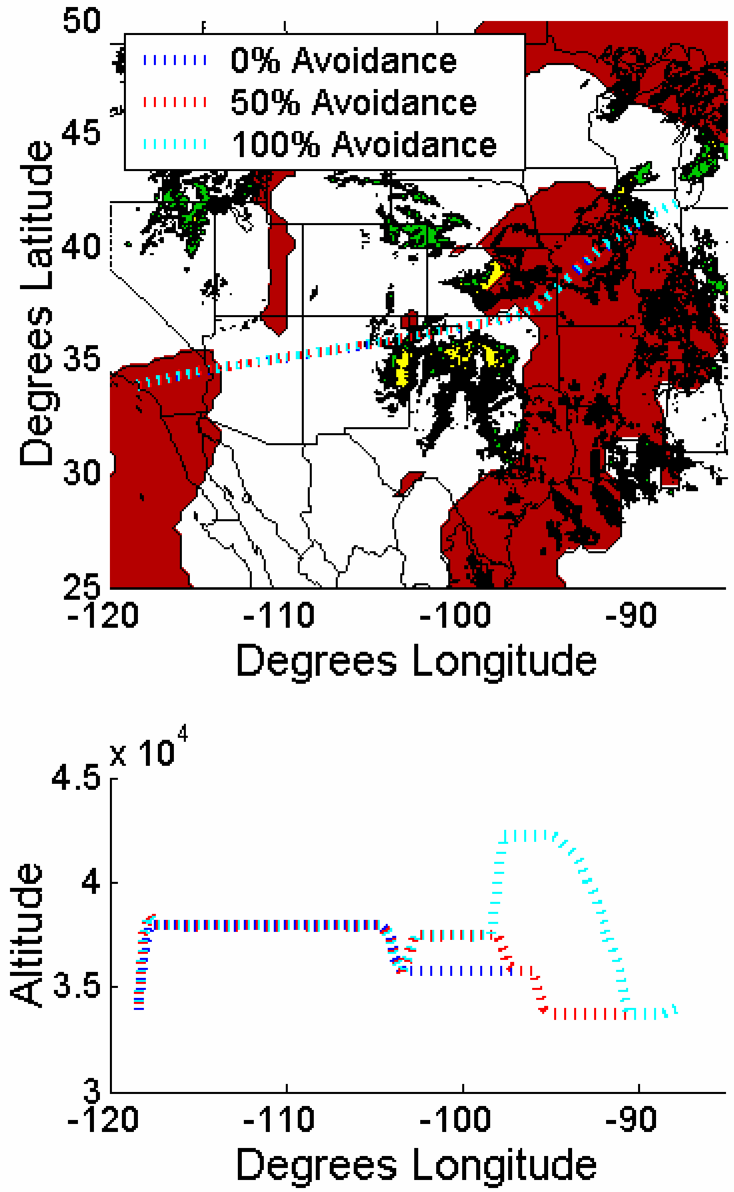

Figure 6.12. Receding horizon trajectories for a scenario involving both persistent contrail mitigation and convective weather avoidance (t = end).

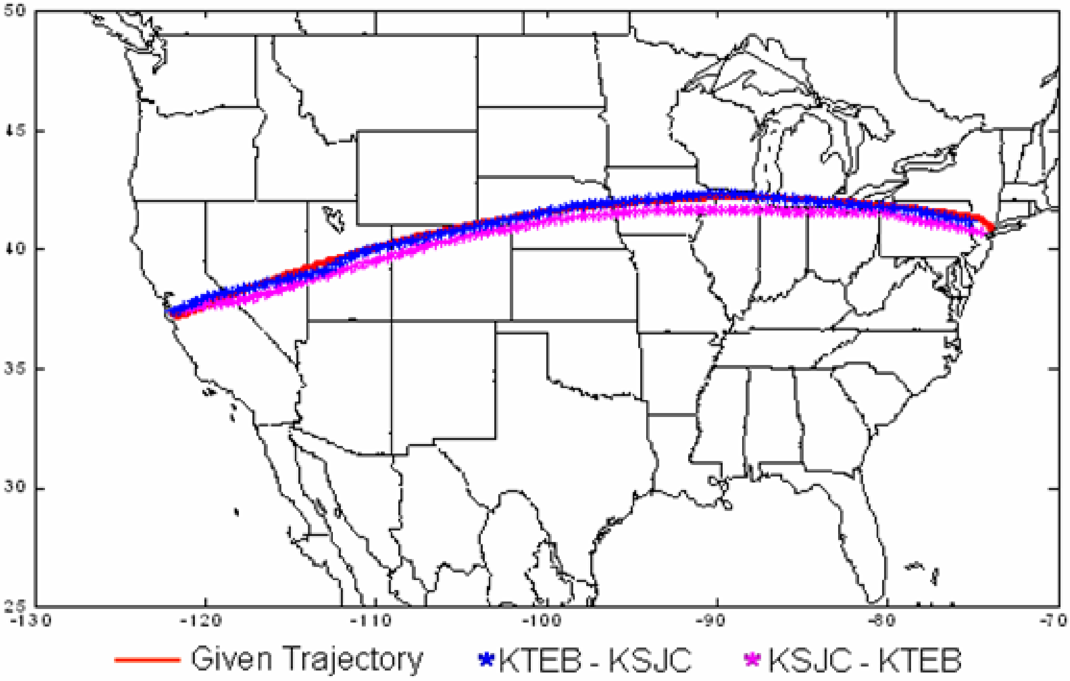

Figure 6.13. Sonic boom impact mitigation trajectories. KTEB is in Teterboro, NJ, and KSJC is in San Jose, CA.

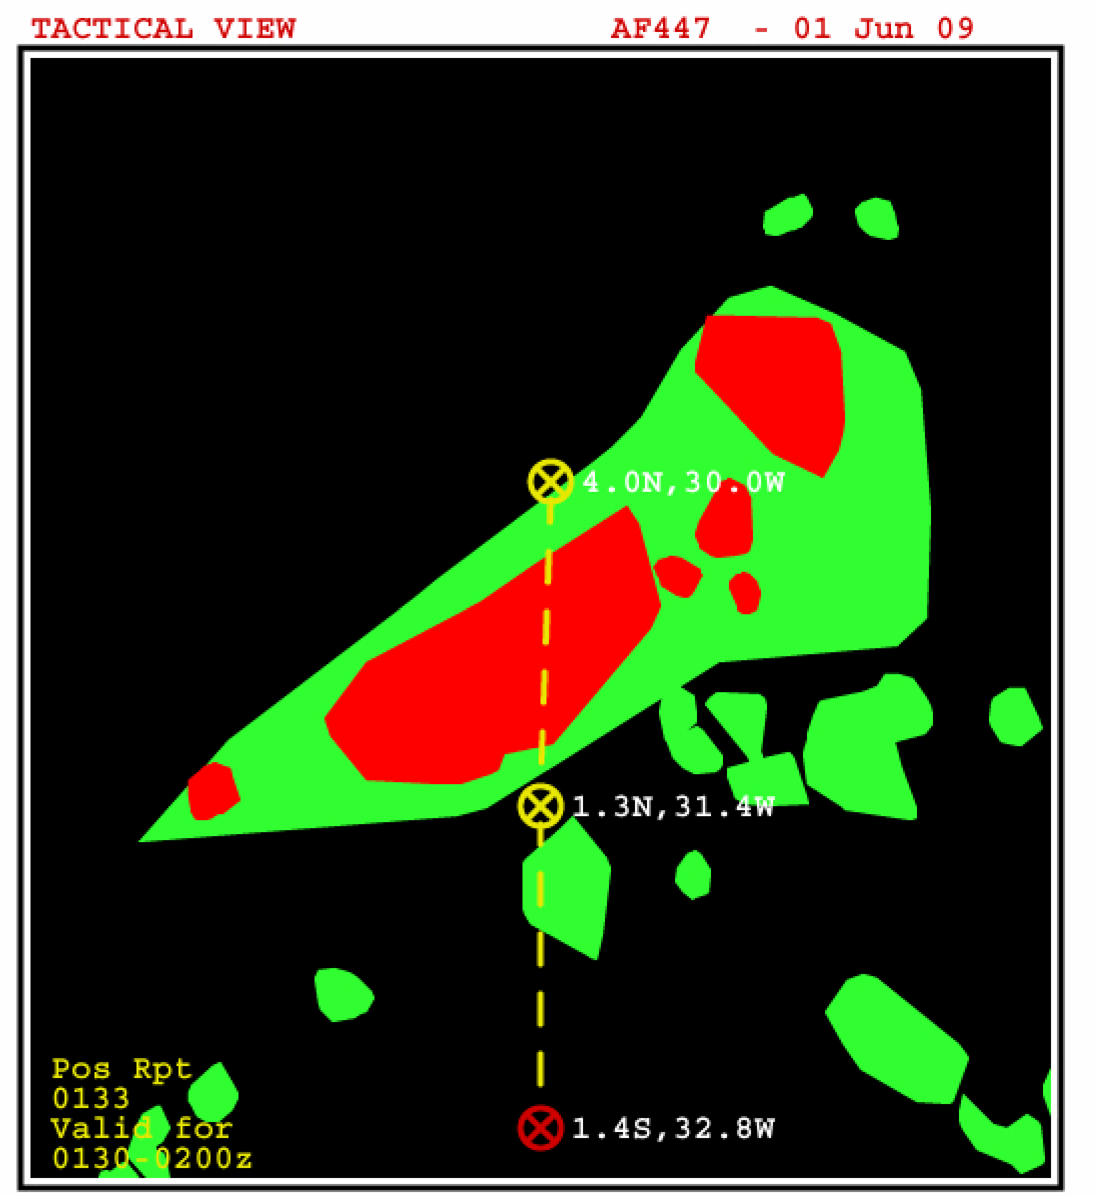

Figure 6.14. Turbulence forecast for June 1, 2009 [74]. The dashed line is the flight path of Air France 447.

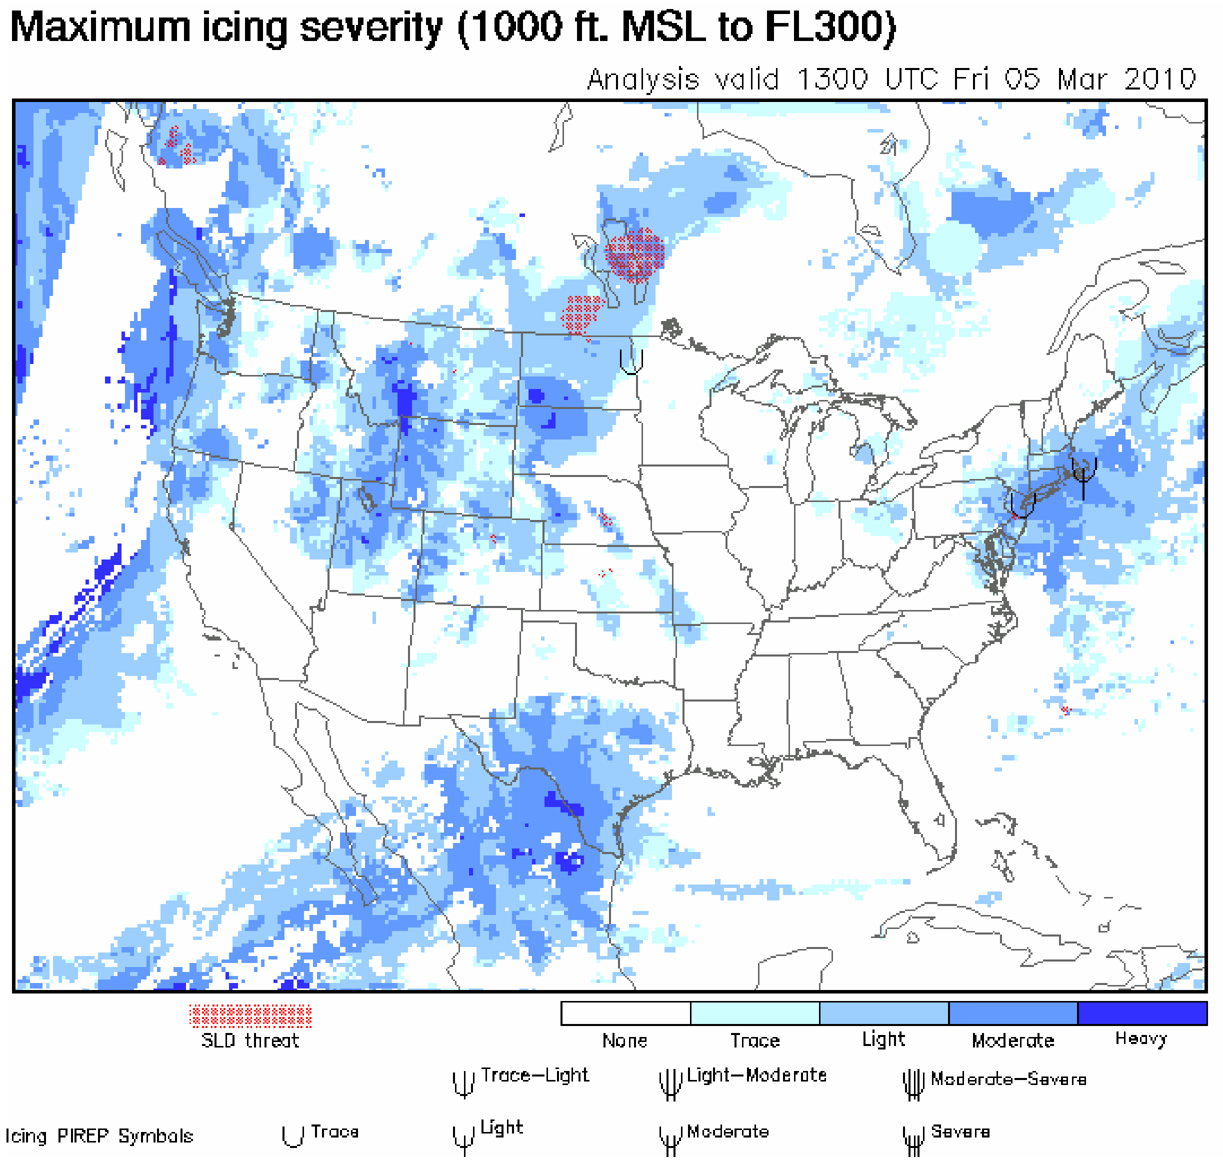

Figure 6.15. Current Icing Product forecast for March 5, 2010.

Chapter 7

Summary, Conclusions, and Recommendations

7.1 Summary

This dissertation presented a path planning framework to deal with trajectory generation problems in the presence of multi-scale obstacles. Chapter 1 introduced path planning in the context of mitigating environmental impact while increasing safety and efficiency in the national air transportation system. It also provided the literature that helped frame the development of the research. Chapter 2 introduced the atmospheric data used to model the environment in the path planning chapters of the dissertation. It also discussed the mechanism of contrail formation, and provided details on the source of the relative humidity data as well as the persistent contrail formation model. The convective weather model was explained, and the source of radar data was discussed. Chapter 3 showed the development of a three-dimensional path planning problem. Much of this path planning algorithm followed from existing work [40, 41]; however, this chapter developed a much more realistic aircraft performance and fuel burn model than previously employed, and analyzed the difference between quadratic and linear cost functions. Chapter 4 presented a novel cost-to-go formulation designed for the inclusion of soft avoidance constraints in the receding horizon MILP. This chapter stepped through the development of the algorithm and provided simulation results for an example scenario. An analysis of the algorithm was conducted to describe how the running time increased with problem size. Chapter 5 gave an analysis of how the receding horizon controller was sensitive to disturbances of multiple scales in the environment. A “best” receding horizon strategy was proposed, which adapted successfully to a problem with multiple scales in the environment. Chapter 6 showed the practical applications of this algorithm. The algorithm was shown to be a viable technique to mitigate persistent contrail formation. The results were compared to existing contrail mitigation strategies. Additionally, the dissertation investigated simulation results for multi-scale scenarios involving both contrail mitigation and thunderstorm avoidance. Lastly, Chapter 6 presented other practical applications for this work including turbulence avoidance, sonic boom mitigation, and aircraft icing avoidance.

7.2 Conclusions

This dissertation presented a framework of a receding horizon path planner for a real time air traffic management application that included an accurate model of aircraft fuel burn, a new receding horizon cost-to-go algorithm, and a model of the atmosphere that included both hard and soft constraints. The performance of the algorithm was tested through simulation, and the practicality of the path planner was investigated though example scenarios using real data. The specific conclusions in the dissertation are listed below:

- This research showed that a more realistic cost function can operate within the bounds of tractability, as long as it is piecewise linear and the number of binary variables in the optimization are bounded. Prior to this dissertation, MILP (mixed-integer linear programming) had mostly been used in relatively simple problems involving the creation of feasible, and time optimal trajectories through a field of obstacles. Recent research in MILP path planning techniques has tended more toward robustness of the algorithm than an assessment of the algorithm’s value in a real-life scenario. Chapter 3 provided a more realistic cost function involving fuel burn as a function of velocity, altitude, acceleration, and climb rate.

- A novel cost-to-go formulation designed for the inclusion of soft avoidance constraints in the receding horizon MILP was developed with this research. Existing research has used receding horizon MILP as a path planning technique for environments containing hard avoidance constraints. The cost-to-go in these formulations restricted the path of least cost in the cost-to-go to pass through the vertices of the obstacles in the environment. This is not an efficient strategy for soft avoidance constraints, where a flight path can pass through the obstacle and incur a penalty. The performance of this algorithm was tested and it was shown to be very close to the performance and shape of trajectories generated with a non-receding horizon MILP formulation.

- The performance of receding horizon MILP was shown to be very sensitive to both the planning horizon and execution horizon lengths in the receding horizon controller. In Chapter 5 the sensitivity of the receding horizon trajectories was investigated in environments containing both hard and soft avoidance constraints that were both static and dynamic. The results of this sensitivity were used to develop a receding horizon strategy with variable planning and execution horizons that adapt to the scale of the obstacles in the environment. The performance of the adaptive strategy was tested through simulation and was shown to perform better than a fixed horizon formulation.

- The path planning algorithm was tested in several real-world examples to show its effect on the environmental impact of aviation. In Chapter 6, the algorithm was applied to the problem of persistent contrail mitigation, and it showed that it can mitigate, on average over a year, 58% of persistent contrails with an increase in fuel consumption of 0.48%. The algorithm was shown to be successful in examples of convective weather avoidance, and in a combined example of persistent contrail mitigation and convective weather avoidance.

7.3 Recommendations

In addition to the contributions listed above, recommendations for future research directions are listed below:

- **Include a more realistic model for the constraints in the National Airspace System (NAS).**

  The research conducted for this dissertation assumed a free-flight type of system for Air Traffic Control (ATC). Operationally, there are many prescribed routes and waypoints that aircraft follow en-route, and there can be unexpected instructions from ATC such as holding and other flow controls. It would be interesting to see how the path planning algorithm in this dissertation would react to such constraints, and how sector capacity would be affected by persistent contrail mitigation and convective weather avoidance.

- **Investigate the effect of parallel computation on the computational performance of the algorithm.**

  As was explained in the dissertation, mixed-integer linear programming is NP-Hard in the number of binary variables. It was shown in the dissertation that receding horizon trajectory accuracy increased with the number of waypoints in the planning horizon; however, larger planning horizons tend to become intractable. Parallel computation should be explored as a possible technique to share the computational burden and allow larger problem formulations to be tractable in real-time.

- **Expand the flight envelope of the aircraft model to include arrival and departure performance and fuel consumption.**

  Arrival and departure routes are greatly affected by convective weather. Furthermore, one of the additional applications listed in Chapter 6 is aircraft icing, which usually occurs during the arrival/approach segment of a flight. Expanding the flight envelope of the aircraft model would enable these scenarios to be more thoroughly investigated.

- **Explore different atmospheric data sources for relative humidity with respect to ice (RHi).**

  The RHi data used in this dissertation were obtained from the Rapid Update Cycle (RUC) model. Although these data are not accurate, they are representative of the size and shape of observed areas of ice super saturation, and were therefore used in this dissertation for the purpose of tool development. In lieu of gridded RHi data, a RHi field could be created using in situ measurements from aircraft en-route. This data field would be less regular than a 3-D grid, therefore presenting an interesting controls problem with limited data.

Bibliography

[1] FAA Air Traffic Organization, “The Economic Impact of Civil Aviation on the U.S. Economy,” 2009.

[2] Ratliff et al., “Aircraft Impacts on Local and Regional Air Quality in the United States,” PARTNER Project 15 Final Report, Oct. 2009.

[3] Maurice L. Q., Lee D. S. (eds) 2009. *Assessing Current Scientific Knowledge, Uncertainties and Gaps in Quantifying Climate Change, Noise and Air Quality Aviation Impacts*. Maurice, L. Q,, Lee, D. S., Wuebbles, D. W., Isaksen, I., Finegold, L., Vallet, M., Pilling, M. and Spengler, J. Final Report of the International Civil Aviation Organization (ICAO) Committee on Aviation and Environmental Protection (CAEP) Workshop, US Federal Aviation Administration and Manchester Metropolitan University, Washington DC and Manchester.

[4] Next Generation Air Transportation System, Federal Aviation Administration Report to the U.S. Congress, 2004.

[5] Intergovernmental Panel on Climate Change, *Aviation and the Global Atmosphere*. J. E. Penner, D. H. Lister, D. J. Griggs, D. J. Dokken, and M. McFarland (eds.), Cambridge University Press. Cambridge, UK, 1999.

[6] Wuebbles, D, et. al. *Workshop on the Impacts of Aviation on Climate Change: A Report of Findings and Recommendations*. Cambridge, MA, June 7-9, 2006.

[7] Travis, D. J., Carleton, A. M., and Lauritsen, R. G., “Contrails reduce daily temperature range,” Nature. Vol 418. 8 August 2002. p. 601.

[8] Waitz, I., Townsend, J., Cutcher-Gershenfeld, J., Greitzer, E., and Kerrebrock, J. Report to the United States Congress: Aviation and the Environment, A National Vision, Framework for Goals and Recommended Actions. Partnership for AiR Transportation Noise and Emissions Reduction, MIT, Cambridge, MA, 2004.

[9] United States Environmental Protection Agency. “Aircraft contrails Factsheet,” EPA430-F-00-005, September 2000.

[10] Brewer, A. W., “Condensation Trails,” *Weather 1*. 34-40, 1946.

[11] Platt, C. M. R., “The effect of cirrus of varying optical depth on the extraterrestrial net radiative flux,” *Q. J. R. Meteorol. Soc*. 107, 671-678, 1981.

[12] Burkhardt, U., Karcher, B., Mannstein, H., and Schumann, U., “Climate impact of contrails and contrail cirrus,” Aviation Climate Change Research Initiative (ACCRI), Jan. 25, 2008.

[13] Cooperative Program for Operational Meteorology, Education and Training (COMET®): http://meted.ucar.edu/nas/index.htm.

[14] Krozel, J., Andre, A., and Smith, S., “The future national airspace system: Design requirements imposed by weather constraints,” in Proc. of the AIAA Guidance, Navigation, and Control Conf., 2003, no. 5769, 2003.

[15] Clarke et al., “Development, Design, and Flight Test Evaluation of a Continuous Descent Approach Procedure for Nighttime Operation at Louisville International Airport,” PARTNER report No. PARTNER-COE-2006-002, Jan. 2006.

[16] Clarke et al., “En Route Traffic Optimization to Reduce Environmental Impact,” PARTNER report No. PARTNER-COE-2008-005, Jul. 2008.

[17] Schumann, U., “Formation, properties, and climatic effects of contrails,” *Comptes Rendus Physique*, 6 (2005), 549-565.

[18] Noppel F, Singh R, Taylor M., “Novel engine concept to suppress contrail and cirrus cloud formation,” International Conference on Transport, Atmosphere and Climate, Oxford, UK, June 2006.

[19] Montgomery CJ, Sarofim AF, Preciado I, Marsh ND, Eddings EG, Bozzelli JW. Experimental and numerical investigation of sootreducing fuel additives. AIAA paper 2005-4472. 41st AIAA/ ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, Tucson, Arizona, July 10-13, 2005.

[20] Gierens K, Kärcher B, Mannstein H, Mayer B. Aerodynamically induced formation of contrails. International Conference on Transport, Atmosphere and Climate, Oxford, UK, June 2006.

[21] Klima, K., “Assessment of a Global Contrail Modeling Method and Operational Strategies for Contrail Mitigation,” M.S. Thesis, MIT, 2005.

[22] Mannstein, H., Spichtinger, P., and Gierens, K., “A Note on How to Avoid Contrail Cirrus,” *Transportation Research Part D*, Vol. 10, No. 5, 2005, pp. 421-426.

[23] Williams, V., and Noland, R. B., “Variability of Contrail Formation Conditions and the Implications for Policies to Reduce the Climate Impacts of Aviation,” *Transportation Research Part D*, Vol. 10, No. 4, 2005, pp. 169-280.

[24] Fichter, C., Marquart, S., Sausen, R., and Lee, D. S., “The Impact of Cruise Altitude on Contrails and Related Radiative Forcing,” *Meteorologische Zeitschrift*, Vol. 14, No. 4, Aug. 2005, pp. 563-572.

[25] Campbell, S., Neogi, N., and Bragg, M., “An Optimal Strategy for Persistent Contrail Avoidance,” AIAA Guidance, Navigation, and Control Conference and Exhibit, Honolulu, Hawaii, Aug. 18-21, 2008.

[26] Campbell, S., Neogi, N., and Bragg, M., “An Operational Strategy for Persistent Contrail Mitigation,” AIAA Aviation Technology, Integration, and Operations Conference, Hilton Head, South Carolina, Sept. 18-21, 2009.

[27] Kuhn, K., “Analysis of Thunderstorm Effects on Aggregated Aircraft Trajectories,” *Journal of Aerospace Computing, Information, and Communication*. Vol. 5, April 2008.

[28] DeLaura, R., and Evans, J., “An Exploratory Study of Modeling Enroute Pilot Convective Storm Flight Deviation Behavior,” 12th Conference on Aviation, Range and Aerospace Meteorology, Atlanta, GA, 2006, pp. 1–15, paper p12.6.

[29] Krozel, J., Lee, C., and Mitchell, J., “Turn-Constrained Route Planning for Avoiding hazardous Weather,” *Air traffic Control Quarterly*, Vol. 14(2), 159-182, 2006.

[30] Pannequin, J., Bayen, A. M., Mitchell, I. M., Chung, H., and Sastry, S., “Multiple aircraft deconflicted path planning with weather avoidance constraints,” in The 2007 AIAA Guidance, Navigation and Control Conference, 2007.

[31] Nguyen, M., Alam, S., tang, J., and Abbass, H., “Dynamic weather avoidance trajectories in a traffic constrained enroute airspace,” ALAR Technical Report Series, TR-ALAR-200712005, 2007.

[32] Frazzoli, E., “Robust Hybrid Control for Autonomous Vehicle Motion Planning,” PhD Dissertation, MIT, 2001.

[33] Milan, M., “Real-Time Optimal Trajectory Generation for Constrained Dynamical Systems,” PhD Dissertation, Cal Tech, 2003.

[34] Schouwenaars, T., “Safe Trajectory Planning for Autonomous Vehicles,” PhD Dissertation, MIT, 2005.

[35] Bryson, A. E., and Ho, Y. C., *Applied Optimal Control*. John Wiley & Sons Inc, 1979.

[36] Hwang, Y. K., and Ahuja, N., “A Potential Field Approach to Path Planning,” *IEEE Trans. on Robotics and Automation*, Vol. 8, No. 1, Feb. 1992.

[37] Ross, I., and Fahroo, F., “Pseudospectral Knotting Methods for Solving Nonsmooth Optimal Control Problems,” AIAA Journal of Guidance, Control, and Dynamics, 27(3):397–405, May-June 2004.

[38] Franz, R., Milam, M., , and Hauser, J., “Applied Receding Horizon Control of the Caltech Ducted Fan.” In Proceedings of the IEEE American Control Conference, 2002.

[39] Milam, M., Mushambi, K., and Murray, R., “A New Computational Approach to Real-Time Trajectory Generation for Constrained Mechanical Systems.” In Proceedings of the IEEE Conference on Decision and Control, pages 845–851, Washington DC, 2000.

[40] Bellingham, J., Richards, A., and How, J. P., "Receding Horizon Control of Autonomous Aerial Vehicles," Proceedings of the IEEE American Control Conference, May 2002, pp. 3741-3746.

[41] Kuwata, Y., and How, J. P., "Three Dimensional Receding Horizon Control for UAVs,” Proceedings of the AIAA Guidance, Navigation, and Control Conference, Aug 2004. AIAA-2004-5144.

[42] Culligan, K., “Nap of the Earth Trajectory Design using MILP,” M.S. Thesis, MIT, 2006.

[43] Schouwenaars, T., Feron, E., de Moor, B., and How, J. P., "Mixed Integer Programming for Multi-vehicle Path Planning, " Proceedings of the European Control Conference, European Union Control Association, Porto, Portugal, September, 2001, pp. 2603-2608.

[44] Ma, C. S., and Miller, R. H., “Mixed Integer Linear Programming Trajectory Generation for Autonomous Nap-of-the-Earth Flight in a Threat Environment,” IEEE Aerospace Conference, 2005.

[45] Chaudhry, A., Misovec, K., and D’Andrea, R., “Low Observability Path Planning for an Unmanned Air Vehicle Using Mixed Integer Linear Programming,” *43rd IEEE Conference on Decision and Control*, Paradise Island, Bahamas, Dec. 14-17, 2004.

[46] Garey, M. R. and Johnson, D. S., *Computers And Intractability: A guide to the Theory of NP-Completeness*, W. H. Freeman and Co., New York, NY, USA, 1979.

[47] Papadimitriou, C. H. and Steiglitz, K., *Combinatorial Optimization: Algorithms and Complexity*, Dover Publications, Inc., Mineola, NY, USA, 1998.

[48] Earl, Matthew. G. D'Andrea, R., “Iterative MILP methods for vehicle-control problems," IEEE Transactions on Robotics and Automation, Vol. 21, No. 6, 2005, pp. 1158-1167.

[49] Vitus, M. P., Pradeep, V., Hoffmann, G. M., Waslander, S. L., and Tomlin, C. J., “Tunnel-MILP: Path Planning with Sequential convex Polytopes,” AIAA Guidance, Navigation, and Control Conference, Honolulu, HI, 2008.

[50] Schouwenaars, T., Valenti, M., Feron, E., and How, J., “Implementation and Flight test Results of MILP-based UAV Guidance,” IEEE Aerospace Conference, Big Sky, MT, Mar. 5-12 2005.

[51] Appleman, “The Formation of Exhaust Condensation Trails by Jet Aircraft”. Bulletin American Meteorological Society, Vol.34, No. 1, January, 1953. pp.14-20.

[52] Schumann, U., “Influence of Propulsion Efficiency on Contrail Formation”. Aersp. Sci. Technol. Vol. 4. 2000. p.391-401.

[53] Gierens, K., Spichtinger, P., “On the size distribution of ice supersaturated regions in the upper troposphere and lowermost stratosphere,” *Ann. Geophys*., 18 (2000) 1687-1690.

[54] Spichtinger, P., Gierens, K., Leiterer, U., and Dier, H., “Ice supersaturation in the tropopause region over Lindenberg, Germany,” *Meteor. Z.*, 12 (2003) 143-156.

[55] Stuber, N., Forster, P., Rädel, G., and Shine, K., “The importance of the diurnal and annual cycle of air traffic for contrail radiative forcing,” *Nature*. Vol. 441, 15 Jun. 2006, pp. 864-867.

[56] Sausen et al., “2005: Aviation radiative forcing in 2000: An update on IPCC (1999),” *Meteor. Z*., 4 (2005) 555-561.

[57] Spictinger, P., Gierens, K., and Read, W., “The global distribution of ice-supersaturated regions as seen by the microwave limb sounder,” *Q. J. R. Meteorol. Soc.* 129, 3391-3410, 2003.

[58] Gettleman, A., Fetzer, E. J., Eldering, A., and Irion, F. W., “The global distribution of saturation in the upper troposphere from the Atmospheric Infrared Sounder,” *J. Clim*. 19, 6089-6103, 2006.

[59] Gierens, K., Schumann, U., Helten, M., Smit, H., and Marenco, A., “A distribution law for relative humidity in the upper troposphere and lower stratosphere derived from three years of MOZAIC measurements,” *Ann. Geophys.* 17, 1218-1226, 1999.

[60] Gibson, J. K., Kallberg, P., Uppala, S., Hernandez, A., Nomura, A., and Serrano, E., “ERA Description,” *ECMWF Re-Analysis Project Report Series*, 1, 1-72, 1997.

[61] Benjamin, S. G., D. Devenyi, S. S. Weygandt, K. J. Brundage, J. M. Brown, G. A. Grell, D. Kim, B. E. Schwartz, T. G. Smirnova, T. L. Smith, and G. S. Manikin, “2004: An hourly assimilation/forecast cycle: The RUC,” Mon. Wea. Rev., 132, 495-518 (Feb. issue).

[62] Weber, M., Evans, J., Wolfson, M., DeLaura, R., Moser, B., Martin, B., Welch, J., Andrews, J., and Bertsimas, D., “Improving air traffic management during thunderstorms,” IEEE Digital Avionics Systems Conference, 2005.

[63] Fahey, T., and Rodenhuis, D., “Continual evolution of CCFP-User needs for extended range prediction,” 11th Conference on Aviation, Range, and Aerospace Meteorology, Hyannis, MA, 2004.

[64] Robinson, M., Evans, J., and Crowe, B., “En route weather depiction benefits of the NEXRAD vertically integrated liquid water product utilized by the corridor integrated weather system,” Aviation, Range and Aerospace Meteorology, ARAM, Portland, OR, 2002, pp. 120–123.

[65] Yoder, T., “Development of Aircraft Fuel Burn Modeling Techniques with Applications to Global Emissions Modeling and Assessment of the Benefits of Reduced Vertical Separation Minimums,” M.S. Thesis, MIT, 2007.

[66] Jardin, M. R., “Ideal Free Flight through Multiple Aircraft Neighboring Optimal Control,” *Proceedings of the American Control Conference*, Chicago, IL, June, 2000, pp. 2879-2885.

[67] Roskam, J., *Airplane Design, Part IV: Preliminary Calculation of Aerodynamic, Thrust and Power Characteristics*, DARcorporation, Lawrence, KS, 2000.

[68] Anderson, J. D., *Aircraft Performance and Design*, McGraw-Hill, 1999.

[69] PERF: Engine Performance Analysis Program v4.2

[70] Lambert, M., Munson, K., Taylor, M, J, H., and Taylor, J. W. R., (eds.) *JANES’S ALL THE WORLD AIRCRAFT*, Jane’s Information Group, Alexandria, VA, 1990.

[71] Padilla, C. E., *Optimizing Jet Transport Efficiency*, McGraw-Hill, 1996.

[72] Maciejowski, J., *Predictive Control with Constraints*. Prentice Hall, 2002.

[73] Mayne, D., Rawlings, J., Rao, C., and Scokaert, P. “Constrained Model Predictive Control: Stability and Optimality.” *Automatica*, 36:789–814, 2000.

[74] Murray, R., and Dunbar, W., “Controller Synthesis for Constrained Flight Systems via Receding Horizon Optimization,” 2nd AIAA “Unmanned Unlimited” Systems, Technologies, and Operations-Aerospace, Land, and Sea Conference, San Diego, CA, Sept.15-18, 2003.

[75] Richards, A., “Trajectory Optimization Using Mixed-Integer Linear Programming,” M.S. Thesis, MIT, 2002.

[76] Cormen, T., Leiserson, C., and Rivest, R., *Introduction to Algorithms*. McGraw-Hill, 1990.

[77] ILOG, *ILOG CPLEX User’s guide*, 2007.

[78] Garcia, C., Prett, D., and Morari, M., “Model predictive control: Theory and practice – a survey,” *Automatica*, 25(3):335–348, 1989.

[79] Stephanopoulos, G., Karsligil, O., and Dyer, M., “Multi-Scale Aspects in Model-Predictive Control,” *Journal of Process Control*, Vol. 10, Apr. 2000, pp. 275-282.

[80] Stuber, N., and Forster, P., “The Impact of Diurnal Variations of Air Traffic on Contrail Radiative Forcing,” *Atmos. Chem. Phys*., vol. 7, 2007, pp. 3153-3162.

[81] Rachami, J., Page, J., Zhao, L., and Kim, B., “Environmental Modeling of Advanced Vehicles in NextGen,” AIAA Aviation Technology, Integration, and Operations Conference, Hilton Head, South Carolina, Sept. 18-21, 2009.

[82] Rachami, J., and Page, J., “Sonic Boom Modeling of Advances Supersonic Business Jets in NextGen,” AIAA Aerospace Sciences Meeting, Orlando, Florida, Jan. 4-7, 2010.

[83] Williams, J., Sharman, R., and Kessinger, C., “Developing a Global Atmospheric Turbulence Decision Support System for Aviation,” Seventh Conference on Artificial Intelligence and its Applications to the Environmental Sciences, Phoenix, AZ, 2009.

[84] NCAR, “New NCAR System May Guide Tranoceanic Flights Around Storms and Turbulence,” http://www.ucar.edu/news/releases/2009/ocean-air-turbulence.jsp#. (accessed 03/12/2010).

[85] NTSB Accident Database & Synopses, http://ntsb.gov/ntsb/query.asp, (accessed 03/15/2010).

[86] Bernstein et al., “Current Icing Potential: Algorithm Description and Comparison with Aircraft Observations,” *Journal of Applied Meteorology*, Vol. 44. Jul. 2005. pp. 969-986.

Author’s Biography

Scot received his B.S. in Aeronautical and Astronautical Engineering from the University of Illinois in 2003. After his undergraduate studies, Scot worked as an aerodynamics engineer for The Boeing Company on the X-45 Joint Unmanned Combat Air System project. In 2005, he returned to the University of Illinois to pursue his M.S. in Aerospace Engineering under the mentoring of Professor Michael Bragg, which he obtained in 2006. Since 2006, Scot has been working on his Ph.D. in Aerospace Engineering, where his research concerns the development of an aircraft path planner to reduce the environmental impact of aviation. In addition to his graduate studies, Scot has been a part-time flight instructor at the University of Illinois Institute of Aviation since 2005. He holds a commercial pilot’s license for single and multi-engine airplanes as well as a flight instructor certificate. Recently, he was awarded the 2010 Roger A. Strehlow Memorial Award for outstanding research accomplishments. This summer Scot will start a new career with MIT Lincoln Laboratory, where he will research decision support tools for air traffic control.