

**WIND MODELS AND STOCHASTIC PROGRAMMING
ALGORITHMS FOR EN ROUTE TRAJECTORY
PREDICTION AND CONTROL**

A Dissertation
Presented to
The Academic Faculty

by

Clayton Tino

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Aerospace Engineering

Georgia Institute of Technology
December 2013

Copyright © 2013 by Clayton Tino

WIND MODELS AND STOCHASTIC PROGRAMMING
ALGORITHMS FOR EN ROUTE TRAJECTORY
PREDICTION AND CONTROL

Approved by:

Professor John-Paul B. Clarke,
Committee Chair, Advisor
School of Aerospace Engineering
Georgia Institute of Technology

Professor Panagiotis Tsiotras
School of Aerospace Engineering
Georgia Institute of Technology

Professor Vitali Volovoi
School of Aerospace Engineering
Georgia Institute of Technology

Professor Eric Feron
School of Aerospace Engineering
Georgia Institute of Technology

Professor David Goldsman
School of Industrial and Systems
Engineering
Georgia Institute of Technology

Date Approved: August 9, 2013

Dedicated to Cindy.

ACKNOWLEDGEMENTS

Upon completion of what initially seemed like an insurmountable task, I would like to acknowledge those in my life who have encouraged and supported my growth both as a scholar and as a person.

This work was made possible due to support provided by the General Electric Company. I would like to thank Mike Durling and Tom Tomlinson at GE Global Research for taking a chance on a young, overconfident intern.

To my advisory committee members, Drs. Tsiotras and Volovoi, your feedback throughout the proposal and defense processes were invaluable in shaping my approach. Your criticisms encouraged me to consider what was beyond my initial limited view of the project and allowed me to adopt a more holistic approach to my work. To my examination committee members Drs. Feron and Goldsman, thank you for taking the time to lend your insight into not only my approach, but the ultimate value of my work as well. I realize that reviewing a dissertation can often be a mundane and thankless job, and I appreciate the enthusiasm that you demonstrated for both the work itself and my ultimate success.

I would like to thank Jim Brooks for keeping me grounded and reminding me that academic pursuit without practical application only serves to boost my own ego.

To my friend and mentor, Liling Ren, I will always appreciate the time you spent with me regardless of whether or not you had any to spare. You may not have realized it at the time, but our afternoon chats were a constant source of inspiration to be a better engineer and person; you set an impossible example that will always be with me.

Any success I have achieved to this point would not be possible without the

guidance of my adviser, John-Paul Clarke. J-P, your mentorship set me on a path that I would have never imagined for myself when I began working for you as a naive undergraduate. More importantly, however, your friendship and constant advocacy of my personal ability and value have done more for me than you will ever know. I look forward to the many years ahead of us, not only as your peer, but as your friend.

Lastly, but most importantly, I would like to thank my mother, Cindy Tino. My earliest memories are of you, Clifford the Big Red Dog, and the joy that learning brought to my life. You have been and always will be my first and favorite teacher, and your love and sacrifice has always been my most significant source of inspiration.

TABLE OF CONTENTS

DEDICATION	iii
ACKNOWLEDGEMENTS	iv
LIST OF TABLES	viii
LIST OF FIGURES	ix
SUMMARY	xi
I INTRODUCTION	1
1.1 Thesis Organization	5
II WIND FORECAST UNCERTAINTY MODELS	7
2.1 Literature Review	8
2.1.1 Data Sets	8
2.1.2 Forecast Uncertainty	14
2.2 Research Gaps	19
2.3 Modeling Methodology	21
2.3.1 A Markovian Approach	22
2.3.2 Data Mappings	27
2.3.3 A Sample Construction	33
2.4 Selected Results	61
2.4.1 Jet Stream Effects	61
2.4.2 Updated Forecasts	65
III THE RTA PROBLEM	69
3.1 Stochastic Programming	69
3.1.1 A Simple Recourse Model	69
3.1.2 Relationship to Dynamic Programming	71
3.1.3 Scenario Trees	74
3.1.4 Nonlinearities	78

3.2	Research Gaps	81
3.3	Algorithm Formulation	82
3.3.1	Scenario Generation	82
3.3.2	Constraints	85
3.3.3	Aircraft Model	89
3.3.4	Search Heuristic	93
3.4	Sample Results	101
3.4.1	Time to Solution	105
3.4.2	Sensitivity to Decision Waypoint	107
3.4.3	Sensitivity to Forecast Uncertainty Conditions	108
IV	CONCLUSIONS	110
4.1	Contributions to the Field	113
4.2	Future Work	114
	APPENDIX A — AIRCRAFT ROUTES	115
	REFERENCES	123

LIST OF TABLES

1	AWIPS 252 grid projection parameters.	29
2	Sample route used to construct uncertainty model for flights between KSEA and KATL.	36
3	Sample route segments used to construct uncertainty model for flights between KSEA and KATL.	38
4	Sample route segments with artificial waypoints for KSEA to KATL scenario.	41
5	Wind forecast uncertainty model analysis routes.	61
6	Notional RTA scenarios along sample routes.	101
7	Notional RTA scenario results.	103
8	Sensitivity to decision waypoint location.	107
9	Route used to construct uncertainty model for flights between KATL and KSEA.	115
10	Route used to construct uncertainty model for flights between KLAX and KATL.	117
11	Route used to construct uncertainty model for flights between KATL and KLAX.	119
12	Route used to construct uncertainty model for flights between KEWR and KATL.	121
13	Route used to construct uncertainty model for flights between KATL and KEWR.	122

LIST OF FIGURES

1	RUC error correlations as a function of observation separation.	14
2	The AWIPS 252 grid.	29
3	RUC-matched ACARS data density in RUC cells along the KSEA to KATL route.	35
4	KSEA to KATL sample route in terms of standard waypoints.	37
5	FAA illustration of fly-by and fly-over waypoints.	39
6	KSEA to KATL sample route in terms of artificial waypoints.	43
7	Filtered ACARS tracks along KSEA to KATL sample route.	47
8	Wind Interpolation scenario.	48
9	Normalized histogram of RUC forecast errors at MWH for the base case model.	52
10	Normalized histogram of RUC forecast errors at A1 for the base case model.	52
11	Cumulative histogram of RUC forecast errors at MWH for the base case model.	53
12	Cumulative histogram of RUC forecast errors at A1 for the base case model.	53
13	Example of state transition accounting.	55
14	Illustration of transition probability matrix ($P^{MWH,A1}$) for the East component of the base case model.	56
15	Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model.	59
16	Illustration of the calculated inhomogeneous Markov chain for the North component of the base case model.	59
17	Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model for the KSEA to KATL route.	63
18	Illustration of the calculated inhomogeneous Markov chain for the North component of the base case model for the KLAX to KATL route.	63
19	Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model for the KATL to KEWR route.	64

20	Illustration of the calculated inhomogeneous Markov chain for the North component of the base case model for the KATL to KEWR route.	64
21	Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model for the KATL to KSEA route. . .	66
22	Illustration of the calculated inhomogeneous Markov chain for the North component of the base case model for the KATL to KLAX route.	66
23	Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model for the KATL to KLAX route. . .	67
24	Illustration of the calculated inhomogeneous Markov chain for the East component of the updated case model for the KATL to KLAX route.	67
25	Sample scenario tree.	75
26	Sample decision scenario for route 1 between KSEA and KATL. . . .	76
27	A two-stage scenario tree for a flight between KSEA and KATL. . . .	85
28	Effects of forecast uncertainty on arrival time distributions versus estimation distance.	87
29	Wind triangle.	89
30	Stage one fuel burn minimization subproblem functional diagram. . .	94
31	Stage one objective calculation functional diagram.	95
32	Stage-two RTA adherence subproblem functional diagram.	96
33	Algorithm time to solution as a function of scenario set size for sample scenario 1.	106
34	KATL to KSEA route with standard and artificial waypoints.	116
35	KLAX to KATL route with standard and artificial waypoints.	118
36	KATL to KLAX route with standard and artificial waypoints.	120
37	KEWR to KATL route with standard and artificial waypoints.	121
38	KATL to KEWR route with standard and artificial waypoints.	122

SUMMARY

There is a need for a fuel-optimal required time of arrival (RTA) mode for aircraft flight management systems capable of enabling controlled time of arrival functionality in the presence of wind speed forecast uncertainty. A computationally tractable two-stage stochastic algorithm utilizing a data-driven, location-specific forecast uncertainty model to generate forecast uncertainty scenarios is proposed as a solution. Three years of Aircraft Communications Addressing and Reporting Systems (ACARS) wind speed reports are used in conjunction with corresponding wind speed forecasts from the Rapid Update Cycle (RUC) forecast product to construct an inhomogeneous Markov model quantifying forecast uncertainty characteristics along specific route through the national airspace system. The forecast uncertainty modeling methodology addresses previously unanswered questions regarding the regional uncertainty characteristics of the RUC model, and realizations of the model demonstrate a clear tendency of the RUC product to be positively biased along routes following the normal contours of the jet stream. A two-stage stochastic algorithm is then developed to calculate the fuel optimal stage one cruise speed given a required time of arrival at a destination waypoint and wind forecast uncertainty scenarios generated using the inhomogeneous Markov model. The algorithm utilizes a quadratic approximation of aircraft fuel flow rate as a function of cruising Mach number to quickly search for the fuel-minimum stage one cruise speed while keeping computational footprint small and ensuring RTA adherence. Compared to standard approaches to the problem utilizing large scale linear programming approximations, the algorithm performs significantly better from a computational complexity standpoint, providing solutions in fractional power time while maintaining computational tractability in on-board systems.

CHAPTER I

INTRODUCTION

The Federal Aviation Administration forecasts that there will be a 50% increase in the number of operations at FAA managed facilities by 2025 [28]. The Next Generation Air Transportation System (NextGen) concept was developed to ensure that air carriers will be able to continue moving people and goods safely and efficiently through the national airspace system (NAS) during this period of growth. The Joint Planning and Development Office (JPDO), tasked with guiding and supporting the creation of NextGen, has stated that trajectory and performance-based operations (TBO) are a primary functional component of the NextGen concept [29]. Primarily, four-dimensional trajectories (4DT) will be the basis for all traffic management functions across all time horizons for managing traffic in high-density and highly complex airspace [27]. The inclusion of 4DT-based technologies is meant to mitigate the impact of increased traffic loads on delays, cost of operation, and the environment by improving both the aircraft's ability to meet schedule constraints mid-flight, as well as the ground's ability to foresee and adjust to operational uncertainties.

A cornerstone of TBO and 4DT is the controlled time of arrival (CTA) functionality. From an aircraft standpoint, a CTA is enabled via use of the required time of arrival (RTA) mode in the aircraft's flight management system (FMS); an RTA mode gives the aircraft the ability to self-deliver itself to an air waypoint at a predefined time. An aircraft's capability to manage its own arrival time to traffic management fixes is necessary to move forward with 4DT, specifically time-based metering applications [29]. Moreover, it has been demonstrated that an aircraft-based solution is preferable to ground-based systems for aircraft speed management in time-based

metering applications due to pilot preference, RTA adherence performance, and air-ground data link challenges [90]. This preference was reinforced by later findings that air-based RTA solutions performed as well or better than ground-based solutions, particularly in the presence of wind forecast error [76, 35].

In the most general sense, an RTA algorithm calculates the speed schedule required to meet an assigned RTA by considering the along-track distance to the RTA fix, a model of aircraft performance capabilities, and the forecasted winds along the aircraft’s planned route. Assuming a constant ground speed (the simplest case), one can view the problem as a distance vs. time calculation. The FMS first generates an estimated time of arrival (ETA) at the RTA waypoint:

$$t_{ETA} = \frac{d_{RTA}}{|\mathbf{v}_g|} \quad (1)$$

where t_{ETA} is the ETA at the RTA waypoint, d_{RTA} is the distance to the RTA waypoint, and \mathbf{v}_g is the vector groundspeed of the aircraft. If there is a difference between the ETA and RTA at the RTA waypoint, the FMS must compute the speed change required such that an updated ETA and the RTA coincide.

In many cases, changing an aircraft’s cruise speed in order to meet an assigned RTA will force the aircraft to operate outside of its optimal performance envelope and introduce operational inefficiencies from a fuel burn standpoint [82]. For example, consider an aircraft that is estimated to arrive late over or at an assigned RTA waypoint, the aircraft will be forced to make up time by increasing its air speed. Though contingent on many factors including flight altitude and operating cost index¹, requiring a higher air speed will usually increase overall fuel consumption [68]. More

¹The aircraft cost index (CI) is a carrier-specific engine operating parameter meant to weigh the overall economic benefit of fuel burn against total flight time [72]. Ranging in values from 0 to 200, a 0 value corresponds to a purely fuel-optimal setting [40]. In the context of a time-constrained, fixed-path RTA operation, flight time is constrained, and a 0 setting is appropriate [68]. It should also be noted, that traditional RTA algorithm implementations may search over a variety of cost indices to find the absolute minimum fuel case [89]. However, the aircraft performance model used in this work relates fuel flow to Mach number alone, and this capability is not required to find the fuel-minimum speed.

preferable methods are available to achieve the RTA from an air traffic management standpoint such as path stretching or shortening to change flight times. However, these options are not specifically available outside of air traffic control (ATC) intervention. Therefore, this research effort will focus on a fixed-path solution specifically applicable to aircraft-based algorithms.

Due to computational power limitations, early RTA algorithm implementations relied on simple search and estimation algorithms to find the optimal operating parameters to meet an assigned RTA [23, 32, 82]. An ETA at the RTA waypoint was generated, and the fuel-minimum cost index and cruise speed required to meet the RTA were determined via a grid search. The problem becomes more complicated, however, when one considers that the groundspeed of an aircraft is not constant. Rather, it is a function of the aircraft’s planned airspeed and the winds encountered along the aircraft’s planned route:

$$\mathbf{v}_g = \mathbf{v}_{TAS} + \mathbf{v}_{wind} \quad (2)$$

where \mathbf{v}_g is the aircraft’s vector ground speed, \mathbf{v}_{TAS} is the aircraft’s true air speed vector (the aircraft’s speed relative the air mass in which it is traveling), and \mathbf{v}_{wind} is the wind vector encountered by the aircraft at each point along the aircraft’s flight path. Thus, the simple distance-time relationship posed in Equation 1 is complicated by the fact that the groundspeed changes as a function of the encountered winds along the route. More exactly, the distance flown by the aircraft over any period of time is represented by the integral equation:

$$d = \int_{t_0}^{t_1} \mathbf{v}_g \cdot \Delta t = \int_{t_0}^{t_1} (\mathbf{v}_{TAS} + \mathbf{v}_{wind}) \Delta t \quad (3)$$

Though the true airspeed \mathbf{v}_{TAS} is planned ahead of time in the aircraft’s flight plan, the winds along the aircraft’s route \mathbf{v}_{wind} must be forecasted or treated as an unknown. Inaccurate forecasted wind speeds lead to inaccurate predictions of the aircraft’s ground speeds, and ultimately the estimated flight time to the RTA fix.

The effects of wind speed forecast errors on trajectory prediction and flight time estimation has been thoroughly discussed in the literature [13, 61, 63, 18]. Early RTA searching algorithms failed to account for wind speed forecast error in ETA calculations, often times failing to deliver the aircraft to the RTA waypoint at the required time. More recent RTA algorithm implementations use a traditional feedback control framework to mitigate the impact flight time estimation error [74, 89]. ETA's at the RTA waypoint are monitored as the flight progresses, and if the difference between the ETA and the RTA grows above a certain tolerance, the RTA algorithm recalculates the required airspeed for the remainder of the route. This approach to aircraft-based control enables extremely accurate delivery times to the RTA fix with delivery error on the order of 7 seconds [89]. A dead-band control approach also allows the algorithm to address the issue of RTA required speed control on fuel burn performance by attempting to use the minimum number of control actions required based on the dead-band gain. However, no credence is given to the quality of the wind forecast information used by the FMS. Again, wind speed forecast uncertainty is handled as a disturbance to be corrected rather than proactively addressed.

More contemporary approaches to managing RTA speed control given wind forecast uncertainty frame the issue as a bounding problem. The aircraft's known speed envelope and a stochastic wind speed and temperature uncertainty model are used to create cones of estimated arrival times representing the earliest and latest possible aircraft arrival time at the RTA waypoint which are then relayed to the pilot. Bounding possible arrival times at an RTA waypoint allows pilots to work in coordination with ATC to better plan required speed changes for RTA operations. However, no new attempt is made algorithmically to minimize the fuel burn required to achieve the assigned RTA given the range of arrival times. Work by De Menorval et. al. suggests a more robust approach to the bounding technique by assigning probability levels to arrival time estimates generated by the FMS [22]. This approach enables

the FMS to make speed control decisions based on which arrival scenarios are more likely to occur given a range of arrival times rather than strictly bounding the possible scenarios that may occur.

The literature suggests a move towards probabilistic and stochastic technologies as the next-step in the evolution of FMS RTA algorithms. However, there is no indication of a major shift in the hardware available on board the aircraft. Any solution hoping to quantify the affects of wind speed forecast uncertainty in fuel-optimal speed change decisions to enable RTA operations will need to be able to operate efficiently given a small hardware footprint. To this end, it is proposed that the RTA speed control problem be reformulated in a two-stage stochastic framework, with a computationally efficient algorithm for solving the fuel burn optimization and RTA adherence subproblems. Further, a data-driven, position-based wind speed forecast uncertainty model will be used to generate scenario sets for the stochastic algorithm.

1.1 Thesis Organization

This research effort includes two distinct yet coupled areas of focus: the wind forecast uncertainty model and the RTA algorithm. Although the modeling of wind forecast uncertainty is an independent endeavour, the RTA algorithm relies directly on the model itself. Thus, the wind forecast uncertainty model is first developed in its entirety in Chapter 2. A review of the wind forecast uncertainty modeling literature is presented, the modeling approach is detailed, a sample model is constructed, and selected results are presented. Development of the RTA algorithm follows separately in Chapter 3. An introduction to stochastic programming is presented along with relevant literature pertaining to the RTA problem, the algorithm is presented in detail, and finally selected results utilizing wind scenarios generated using the forecast uncertainty model developed in Chapter 2 are presented. The conclusions from this work are presented in Chapter 4 along with summaries of both the included results

as well as the contribution of this work to the field given the research gaps identified in prior chapters.

CHAPTER II

WIND FORECAST UNCERTAINTY MODELS

The proposed solution involves recasting the RTA speed control problem as a two-stage stochastic program, with a position-based forecast uncertainty model forming the basis of the scenario sets for the stochastic formulation. In the general sense, an uncertainty model is a mathematical model describing the randomness associated with a system. As has been previously discussed, forecast wind speed values along an aircraft's flight path tend to differ from the actual wind speeds experienced by the aircraft as it traverses its planned route. These differences in observed wind speeds versus forecast wind speeds are referred to as forecast "errors." More generally, forecast error is the observable difference between a particular forecast and observation data for the valid forecast time. Forecast error differs from forecast uncertainty, which is an estimate of future forecast error made *a priori*. The forecast uncertainty model developed in this work is therefore a mathematical model that estimates the future random differences between forecast wind speeds along an aircraft's planned route, and what the aircraft ultimately experiences en route.

Before detailing the forecast uncertainty modeling methodology, the relevant data sets in the field will be examined: the Rapid Update Cycle (RUC) forecast product and the Aircraft Communications Addressing and Reporting System (ACARS) data set. A review of the relevant literature concerning wind speed forecast uncertainty modeling and the affects of wind speed forecast uncertainty on aircraft trajectory prediction is then presented. The gaps in the current research efforts to date will then be identified, and the scope of the forecast uncertainty modeling effort will be defined. Finally, the uncertainty modeling technique will be developed in its entirety,

and selected results discussed.

2.1 Literature Review

2.1.1 Data Sets

2.1.1.1 The Rapid Update Cycle

Though most air carriers operate internal weather clearinghouses, many utilize the National Oceanic and Atmospheric Administration’s (NOAA) Rapid Update Cycle Forecast (RUC) product in some capacity. Thus, the RUC forecast product was chosen as the basis of the uncertainty modeling effort. In discussing the RUC, the ACARS data set will be mentioned as well; specific details regarding the ACARS data set are available in Section 2.1.1.2.

The Rapid Update Cycle is an operational mesoscale data assimilation and numerical forecast system (a full specification is available from NOAA). The key aspects of the RUC forecast product differentiating it from other numerical weather prediction systems are its hourly forward assimilation cycle [7], and its use of a hybrid isentropic-terrain-following vertical coordinate for both the assimilation and forecast model components [8], the former being of particular interest to this work. The RUC utilizes a one hour forward intermittent data assimilation cycle to add new observations from various data sources to the current forecast model, using the previous one hour RUC forecast as a basis for the new atmospheric model. More specifically: new observations are compared to the previous one hour forecast to generate forecast residuals; these residuals are then analyzed by the model to produce an estimate of the forecast error field; this error field is then summed with the previous one hour forecast to produce a corrected one hour model. This scheme essentially uses the RUC forecast model in a simplified Kalman filtering framework to introduce new observations to the next hour forecast. Historically, the RUC has roots in the RUC1 model which operated on a three hour assimilation cycle [6]. However, the data available

from observable sources such as rawinsondes, GPS, ship reports, surface stations, and particularly aircraft measurements are available on an hourly cycle, meaning two-thirds of the available observation data was lost. As a result, it was desirable to introduce a higher-frequency forward assimilation framework. The technical details of this framework will not be discussed as they are outside the scope of this work. However, the use of aircraft reports as a high-frequency data source will be examined as they are a primary component of the uncertainty modeling effort.

Aircraft wind speed observations play a primary role in the current RUC framework, as they have been shown to contribute significantly to forecast validity in even the earliest data assimilation forecasts due to both the quality of their reports as well as the volume of data available [3, 4]. However, as the distribution of aircraft reports in both space and time is strictly dependent on the route structure of air carriers, there is a trade-off to be made between the length of the assimilation window used in the model and the horizontal resolution that results due to that choice. Aircraft reports span a shorter distance in space given a shorter observation window meaning there is less geographic coverage as the length of the assimilation window is reduced; this reduction in spatial coverage directly impacts the horizontal resolution of the model. A further temporal error is introduced by the data assimilation method when one considers the time of the aircraft observation to be valid at the time of analysis, though observations rarely occur at the exact time of analysis. Historically, this spatiotemporal error was accepted as a necessary evil of the data assimilation framework. Observations were grouped into time windows several hours wide centered on the analysis time of a particular forecast and considered to be valid [55, 84, 52]; these observations were then used according to their particular assimilation scheme as if they occurred at the time of analysis. However, as the time resolution of the assimilation is increased, the error introduced by time validity of the data observations

becomes more significant. The simplest solution to the time validity issue is to increase the analysis frequency; more frequent analyses corresponds to a lower average time difference between the observation time and the analysis time. The expected consequence of this solution is a decrease in the volume of data available for a particular analysis, thus decreasing the horizontal resolution of the model. One must also ensure that as the assimilation window is shortened, the data density remains sufficient to ensure there are enough valid observations to correct the background model given the desired horizontal resolution of the model [60]. Benjamin et al. examined this phenomenon and demonstrated that the impact of the analysis time window on spatiotemporal validity of the RUC model depends specifically on the feature being modeled. A fast moving feature such as a jet streak (on the order of 30 m/s) was shown to be subject to a 100 km stretching distortion with an assimilation window as short as an hour [7]. In order to ensure the assimilation window was sufficiently short to reduce time validity errors while remaining sufficiently large to include enough aircraft observations to model prevalent propagating features in the atmosphere, the middle-ground position of a one hour assimilation cycle is currently used in the RUC model. However, improvements to the quality and quantity of available observation data, as well as improvements to computational capabilities have led to increases in the horizontal resolution of the RUC along its development cycle. The 1994 release of the RUC1 provided reports at a 60 km horizontal resolution and 25 vertical levels; this work will make use of the RUC20 product, released in 2002, which provides reports at a 20 km horizontal resolution and 50 vertical levels.

Several researchers have explored the accuracy of the RUC versus both archived and live aircraft wind speed reports. Cole et al. performed the first comprehensive review of the RUC model in 1998 [15], in which aircraft wind speed reports from the Meteorological Data Collection and Reporting System (MDCRS) from a 1300 km wide region encompassing the Denver center airspace were compared to the RUC1

forecast over a one year period. The RUC1 was found to exhibit an RMS vector error of 6.74 m/s , with a median error vector of 4.99 m/s . The 90th percentile error vector was shown to be 10.18 m/s with roughly 11% of the RUC error vectors being larger than 10 m/s . For air traffic management purposes, peak errors of greater than approximately 7.7 m/s were found to be detrimental to trajectory prediction purposes as the aggregate position error over a 20 minute prediction becomes larger than the 5 nautical mile minimum en-route radar separation distance [33, 69]. Of greater importance, however, were the correlation amongst errors within the region. For level flights of 20 minutes at 400 kts , the RUC1 demonstrated a linear correlation coefficient of approximately 0.45. Uncorrelated wind speed errors will tend to change sign along the route, ultimately canceling the effect of the error on flight time prediction over the course of the flight. However, correlated error values tend to compound not only as a function of the spatial correlation term, but also the reference time along which they are considered; the longer the flight time, the more significant the effect of correlated errors on flight time prediction [16, 17].

Schwartz et al. performed a later study of the RUC2 40 km resolution grid [78]. Wind speed reports from the ACARS data set were compared to RUC2 forecasts for the same time period and airspace as the Cole study [15]. The RUC2 model showed slight improvement of RMS error vectors, with daily mean RMS values improving roughly 10%. This improvement, though small, was considered to be statistically significant given the large number of ACARS reports used as well as the fact that ACARS observations are unbiased with respect to the forecast error of the RUC model. A more interesting result, however, was that the distributions of errors between the RUC1 and RUC2 models were found to be independent, implying that as model improvements are made, the uncertainty characteristics of the RUC model will change. The Schwartz study further found that seasonal affects exist between summer and winter months with the seasonal movement of the jet stream, with higher RMS

error vectors during the winter months.

2.1.1.2 The Aircraft Communications, Addressing, and Reporting System

United States air carriers have implemented a fully automated meteorological reporting system via a communications system called the Aircraft Communications, Addressing, and Reporting System (ACARS). Relying primarily on air-ground VHF radio links, aircraft are able to report meteorological conditions to ground stations that cover a majority of the U.S. airspace (though newer aircraft are able to make use of satellite communication links). Key parameters of these observation include:

- Longitude and latitude (tenth of a minute).
- Time (nearest minute).
- Temperature (nearest tenth of a degree).
- Flight level (pressure altitude to the nearest hundred feet).
- Wind direction (nearest degree).
- Wind speed (nearest knot).

As introduced in Section 2.1.1.1, these meteorological reports play a key role in both the ongoing development and validation of numerical forecast systems such as the RUC. Accordingly, the accuracy and best use practices of these reports are of key interest when considering ACARS reports for any sort of modeling effort.

Initial studies of wind speed measurements from ACARS equipped aircraft compared observations from collocated aircraft to rawinsonde measurements in order to provide an upper bound limit on ACARS measurement variability [77]. Observations from ascending and descending aircraft separated by less than 150 *km* and 90 minutes in time were shown to exhibit wind speed observational variation of 5.76 *m/s* across 4440 matched data pairs. This variability was further reduced to 4 *m/s* when only

data pairs separated by less than 25 *km* and 15 *min* were considered, and reduced further still to 2.84 *m/s* when only considering the “best” sample set. It should be noted by the reader, however, that this measured error reflects contributions from three known sources: rawinsonde measurement and reporting error, aircraft measurement and reporting error, and mesoscale variability in wind speeds. Of key interest to this work, is the contribution due to aircraft measurement and reporting error (i.e., the error specifically inherent to the ACARS observations themselves). ACARS equipped aircraft compute wind speed observations by comparing the motion vector of the aircraft with respect to the earth as given by the inertial navigation system (INS), to the motion vector of the aircraft with respect to the air as given by the total airspeed and heading measurements. This methodology has been demonstrated to provide high accuracy measurements of wind speed on the order of 0.5 *m/s* [66]. INS wind speed measurements are also heavily influenced by aircraft maneuvers, however. Bisiaux et al. demonstrated significant measurement errors for maneuvers in which the roll angle exceeded 5° [11]. At the time of the 1995 Schwartz study discussed above [78], the ACARS data set contained no indicator of the aircraft’s roll state. Nonetheless, the authors concluded that error due to ACARS wind speed calculations contributed far less to total error calculation than both rawinsonde and mesoscale affects, and that ACARS wind speed errors were likely on the order of the INS limits posed by Nicholls. This conclusion was later confirmed during the RUC verification studies discussed in Section 2.1.1.1, by which time an aircraft roll indication flag had been added to the ACARS standard. In the 1999 verification studies, Benjamin et al. found estimated ACARS wind speed report errors on the order of 1.1 *m/s* [9].

2.1.2 Forecast Uncertainty

By virtue of the data assimilation modeling methodology used to generate RUC forecasts, the forecast itself will never exactly match observation values. The tradeoff between spatial resolution and analysis time discussed in Section 2.1.1.1 results in a model that is unable to represent phenomena of every length and time scale present in nature. This behavior is evident in an examination of the error correlation results described by Cole et al. [15]. The Cole study calculated the correlation between RUC errors as a function of observation separation distance then fit an exponential function to the resultant data. The specific correlation values from this study are not available. However, the exponential fits are presented in Figure 1 (the results are presented in terms of the u and v directional components unique to the RUC). Of importance in Figure 1 are the correlation values as the separation distance ap-

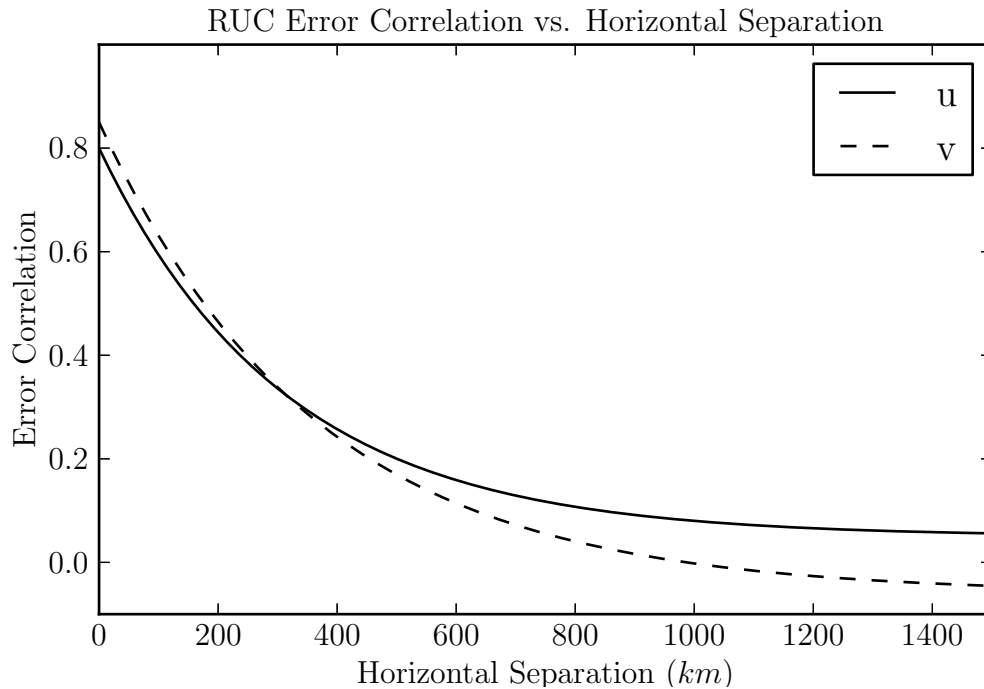


Figure 1: RUC error correlations as a function of observation separation.

proaches zero. By definition, one would expect the correlation between collocated errors to approach unity. This discrepancy is the result of smaller scale atmospheric

features not being captured by the model. The observation network, in this case a series of aircraft reports, is not of sufficient horizontal resolution to describe smaller scale features. Mondoloni provides a description of the phenomenon as follows [61]: consider a network of K observation stations with multiple observations O and reference forecast values F . The correlation between observation errors at two stations ℓ and m can be calculated:

$$\rho_{\ell,m} = \frac{\overline{(O_\ell - F_\ell)(O_m - F_m)}}{\sqrt{\overline{(O_\ell - F_\ell)^2} \overline{(O_m - F_m)^2}}}$$

where the observation values O_ℓ and O_m include any sensor error. A third variable T is now introduced to represent the “true” value of the observation at a station k . It should be noted that T only represents the “true” value of those components of the time and distance spectrums able to be captured by the RUC. Observation and forecast error terms can now be found:

$$E_O^2 = \frac{1}{K} \sum_{k=1}^K \overline{(O_k - T_k)^2}$$

$$E_F^2 = \frac{1}{K} \sum_{k=1}^K \overline{(F_k - T_k)^2}$$

These error terms necessarily include error due to the inability of the RUC to represent smaller scale phenomenon together with measurement errors. The correlation function is now expanded:

$$\rho_{\ell,m} = \frac{\overline{[(O_\ell - T_\ell) - (F_\ell - T_\ell)][(O_m - T_m) - (F_m - T_m)]}}{\sqrt{\overline{[(O_\ell - T_\ell) - (F_\ell - T_\ell)]^2} \overline{[(O_m - T_m) - (F_m - T_m)]^2}}}$$

If one further assumes that the forecast errors are uncorrelated with measurement errors, the correlation function can be further simplified:

$$\rho_{\ell,m} = \frac{\overline{(O_\ell - T_\ell)(F_\ell - T_\ell)} + \overline{(O_m - T_m)(F_m - T_m)}}{E_O^2 + E_F^2}$$

If the errors are considered to be homogeneous, and the station measurement errors uncorrelated in space, the limit of the correlation as the distance r between stations

the station ℓ and m reduces:

$$\lim_{r \rightarrow 0} \rho(r) = \frac{E_F^2}{E_O^2 + E_F^2} \quad (4)$$

The result found in Equation 4 demonstrates the paradox that results by neglecting to account for smaller scale effects present in the atmosphere; this error is referred to as error of “under-represented scales” or “representativeness.”

Three sources of error in the RUC forecast product have been presented to this point:

- Sensor error in predominantly aircraft-based atmospheric observations on the order of 1 *m/s*.
- Error introduced by the trade between assimilation window length, data density, and observation time (discussed in Section 2.1.1.1).
- Error of representativeness due to the density of the observation grid (the reader should note that error of representativeness is coupled with error due to the assimilation scheme, as the assimilation window length will determine the possible density of the observation grid when using aircraft observations).

Several techniques have been developed to minimize the effects of model scale and representativeness, a prominent approach being the time-lagged ensemble forecast. A time-lagged ensemble forecast is simply a weighted mean of time-lagged deterministic weather forecasts valid for the same time. Posed originally by Hoffman and Kalnay in 1983 [36], ensemble forecasts address both error due to the assimilation scheme and error of representativeness by incorporating observations from multiple analysis time windows. Modeling errors introduced by the data assimilation scheme tend to be blunted by the averaging of multiple forecasts, and the density of the effective observation grid is increased as more forecasts are added. Multiple weighting schemes have been introduced, two of which have been explored specifically with the RUC

forecast product. Lu et al. examined the improvement in time-lagged ensemble RUC forecasts for both a simple arithmetic mean and optimized weighting scheme using a linear regression model to minimize the forecast error as compared to a reference “truth” [57]. This study found that both methodologies were effective in their ability to reduce random model error due to data assimilation. However, the regression technique was more capable by virtue of giving higher preference to those forecasts which performed better when compared to observation data.

Though forecast ensembles provide a means of diminishing the impact of forecast errors of scale and representativeness, the question of the forecast uncertainty remains unanswered. The simplest approach to modeling forecast uncertainty is to simply express forecast error in terms of root mean square statistics for the forecast product being used, an approach common in aircraft conflict prediction [49]. However, root mean square statistics fail to completely quantify the uncertainty characteristics of a forecast models like the RUC, as they neglect the spatial correlation of forecast errors due to under-represented scales [34]. Admittedly, a model of the forecast error spatial correlation is the prohibiting factor in many of the research efforts to date:

“[A]lthough known to be a critical issue in aircraft conflict prediction, the wind spatial correlation [of forecast errors] is largely ignored in the current literature, probably due to the difficulty in its modeling and analysis.” [39]

As a result, a standard approach to modeling forecast uncertainty is to assume some stochastic distribution of future forecast errors, usually in the form of a Gaussian or Brownian random field with various correlation structures. Often times the choice in this correlation structure is dependent entirely on the end-use of the uncertainty model. In aircraft conflict detection, for example, the geometry of an individual conflict scenario may determine the correlation parameters used for the uncertainty model [39]. Other times, the choice in correlation structure is based on the desired complexity of the model. In the simplest case, forecast errors may be completely

uncorrelated (i.e., truly random) [12, 24]. In more complex implementations, forecast errors can be correlated in space alone [39], or in both space and time [58, 14, 88, 59]. Interestingly, regardless of the particular correlation scheme chosen, the authors of these works tune their models to produce the same RMS error statistics produced by the Cole’s study of the RUC forecast.

A more complete forecast uncertainty model was proposed by Mondoloni in 2006 in order to address the multiple sources of forecast error [62]. Mondoloni’s solution describes forecast uncertainty as a sum of three random components, each characterizing a different type of forecast error:

$$w_{error} = \varepsilon_R + \varepsilon_S + \varepsilon_L$$

where w_{error} is the total wind speed uncertainty, ε_R is the error due to representativeness of the model, ε_S is the prediction error in those features capable of being modeled at scale, and ε_L is a large-scale constant model bias. The model continues by further assuming a correlation structure for each component: the error due to underrepresented scales is modeled as a zero-mean process with a time correlation function based on the density of the observation network; the prediction error at modeled scales is modeled as another zero mean process with a correlation structure based on work by Hollingsworth and Lönnberg [37]; and the large scale error is modeled as a constant sampled from a normal distribution with a specified variance. Via Monte Carlo simulation, Mondoloni demonstrated that this decomposition approach provides forecast error distributions similar to those described by Schwartz et al. [78].

Though decomposing forecast uncertainty into separate processes based on the error scale provides results more consistent with the RUC validation studies, it fails to address the issue of localization. A more recent study by Lee et. al compared flight time predictions along paths through ensembled RUC forecasts to examine the affects of RUC uncertainty on trajectory prediction error. Though not the focus of the study, Lee demonstrated that different paths through the airspace exhibit different

uncertainty characteristics [51]. Moreover, the RUC forecast exhibits highly regional uncertainty characteristics that have not yet been captured. The majority of works to date rely on the results of the Cole and Schwartz RUC verification studies as the basis for any wind speed forecast uncertainty model. However, tuning an uncertainty model to fit the RMS statistics for one region is not acceptable for every application, particularly long-scale trajectory prediction in which an aircraft will traverse several unique regions of the RUC model.

Recent work by Zheng and Zhao has attempted to address the issue of regional variations by using weighted ensembles of forecast error statistics [91]. Zheng and Zhao propose a non-Gaussian random field composed of weighted sums of historical forecast errors at certain “look ahead” distances and two purely random components to represent measurement errors in the nominal wind speed as well as modeling errors. Again, however, the authors assume that forecast uncertainty correlation is represented by an exponential curve as discussed in reference to Figure 1. Specifically, forecast error records falling within the characteristic correlation length (found to be between 250 *m* to 350 *m*) are summed with the weights being inversely proportional to the distance from the local point squared. As a result, the model is inherently non-stationary as one sums the contributions of neighboring measurements while moving along a specified path.

2.2 Research Gaps

Though the error characteristics of the RUC forecast have been studied extensively, an uncertainty model capturing both errors in model scale and representativeness as well as the unique geographical characteristic of the RUC model is still lacking. RMS statistics and Gaussian random field models are completely appropriate for short range trajectory predictions such as those used for conflict detection, but are not satisfactory when considering long scale trajectory predictions through several

regions of the national airspace system. There is a clear need for a forecast uncertainty model that not only addresses the issues of model scale and representativeness, but also captures the unique regional uncertainty characteristics of the RUC forecast.

2.3 Modeling Methodology

Given the need for a position-specific forecast uncertainty model as identified in Section 2.2, focus must now shift to defining a methodology that produces an uncertainty model that captures the effects of multiple model scales including error of representativeness while also adequately capturing the regional error correlation characteristics of RUC forecasts. Again, consider the end-use scenario for forecast uncertainty model, an aircraft in cruise: there exists a flight plan detailing the aircraft’s future planned path, a RUC wind speed forecast along this route, and a measure of the current wind speed at the aircraft’s current and past positions. The forecast error vector at any point along the flight path can be defined as:

$$\boldsymbol{\varepsilon} = \boldsymbol{w}_{forecast} - \boldsymbol{w}_{actual}, \quad (5)$$

if the sensor error is assumed to be small in comparison to the magnitude of the forecast wind speeds. Now consider the contents of the RUC and ACARS data sets: the RUC provides wind speed forecasts for the cruise regime, and the ACARS data set provides historical aircraft flight paths through the national airspace system (NAS) with corresponding wind speed measurements at selected points along those paths. One can construct a series of forecast error measurements $\boldsymbol{\varepsilon}$ by comparing RUC forecast values $\boldsymbol{w}_{forecast}$ with corresponding ACARS wind speed measurements \boldsymbol{w}_{actual} at different points within the NAS, effectively recreating the data set used in the RUC verification studies but on scale much larger than the Denver TRACON area [15, 78]. This technique was used by the author in 2009 to examine the regional error characteristics of the RUC forecast on a NAS-wide scale [86]. Though this technique was successful in identifying a tendency of the RUC to overemphasize the contribution of the jet stream in certain regions, a major shortcoming of this technique was again a lack of any understanding regarding the correlation between error measurements as a function of position within the NAS. As the issue of error correlation remains

a path-specific problem, any analysis must consider specific paths through regions of the NAS rather than simply the aggregate statistics of those regions.

Conveniently, air carriers tend to follow very similar routes through the NAS depending on the origin and destination airport pair of a given flight, meaning ACARS flight data tends to be clustered along very similar ground tracks when considering a specific origin-destination pair. By approaching the error analysis from a route-based standpoint, it is possible to generate a series of wind speed measurement strings along very similar ground tracks. These measurement strings can then be compared to the corresponding RUC forecast to produce a series of wind speed errors linked to the origin-destination pair. More importantly though, as each ACARS flight is considered in totality rather than as contributing independent wind speed measurements, there are now measurements of prior and future forecast errors at different points along a common route for individual flights. Given a sufficiently large data set, one can construct an uncertainty model that leverages this time history of error transitions as a function of position; a Markov chain is well-suited to this task.

2.3.1 A Markovian Approach

Let S be a countable set, with each $s \in S$ representing a state. A state s is a discrete interval on the range $R = [-a, a]$ where R contains q individual states s . For example, if $R = [-5, 5]$ with $q = 5$ states, S contains the state intervals $\{(-5, -3], (-3, -1], (-1, 1], (1, 3], (3, 5]\}$. In the context of the forecast uncertainty model, each forecast error value calculated by comparing a RUC forecast wind speed with an archived ACARS measurement can be mapped to a state interval s . Using the prior example of $R = [-5, 5]$ with $q = 5$ states, an error value $\varepsilon = -0.5$ maps to state $s_3 = (-1, 1]$. It is important to note that a state s refers to a mapped bin in S representing a range of forecast error values and *not a position along an origin-destination route*. Continuing, $\lambda = (\lambda_s : s \in S)$ is considered to be a measure on S if

$0 \leq \lambda_s \leq \infty$; furthermore, if $\sum_{s \in S} \lambda_s = 1$, λ is a distribution of S .

Now consider a random variable X in a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where X holds values in S as a function $X : S \rightarrow \Omega$. If:

$$\lambda_s = \mathbb{P}(X = s) = \mathbb{P}(\{w : X(w) = s\})$$

then λ defines the distribution of X . That is to say, X models forecast error as a random state which takes value s with a probability λ_s .

A matrix $P = (p_{i,j} : i, j \in S)$ is said to be stochastic if every row $(p_{i,j} : j \in S)$ is a distribution. A Markov chain $(X_n)_{n \geq 0}$ is now defined as a stochastic process holding the following properties in terms of a stochastic transition probability matrix P and initial distribution λ :

1. $\mathbb{P}(X_0 = s_0) = \lambda_{s,0}$
2. $\mathbb{P}(X_{n+1} = s_{n+1} | X_0 = s_0, \dots, X_n = s_n) = p_{s_n, s_{n+1}}$

for $n \geq 0$ and every $s_0, \dots, s_n \in S$. $(X_n)_{0 \leq n \leq N}$ is *Markov* (λ, P) if $(X_n)_{0 \leq n \leq N}$ is a sequence of random variables satisfying (1) and (2) above for $n = 0, 1, \dots, N - 1$. In other words, a Markov process has (1) an initial state distribution $X_0 = \lambda$ and (2) for $n \geq 0$, X_{n+1} has distribution $(p_{i,j} : j \in S)$ conditional on $X_n = s$ and independent of X_0, \dots, X_{n-1} . Formally, we can state that a Markov chain holds the Markov property:

Theorem 2.3.1 *Let $(X_n)_{n \geq 0}$ be Markov (λ, P) . Then, conditional on $X_m = s$, $(X_{m+n})_{n \geq 0}$ is Markov (δ_i, P) and is independent of the random variables X_0, \dots, X_m .*

In the simplest terms, a Markov chain is defined by a set of states $S = \{s_1, s_2, \dots, s_n\}$. The chain starts in one of these states and moves successively between these states over a series of steps. If the chain is in state s_i at step n , then it moves to state s_j at step $n + 1$ with probability $p_{i,j}$. The Markov property (Theorem 2.3.1) holds

that $p_{i,j}$ is independent of the chain's state at any prior step. Howard illustrates a Markov chain as a frog jumping between a series of lilly pads; the frog begins on one pad, then proceeds through a series of neighboring pads based solely on the pad he currently occupies [38].

Recall that distributions and measures λ are merely row vectors, indexed by a state $s \in S$, and the transition probability matrix P is simply a matrix indexed by indices on $S \times S$. Given S is finite, the states s_i can be indexed $1, 2, \dots, N$, meaning λ is simply a $1 \times N$ vector and P an $N \times N$ matrix. Matrix multiplication can be extended to define new measures λP and and a new matrix P^2 [67]:

Definition $(\lambda P)_j = \sum_{i \in S} \lambda p_{i,j}$

Definition $(P^2)_{i,k} = \sum_{j \in S} p_{i,j} p_{j,k}$

where P^n follows similarly for any n . Using these definitions, P^0 refers to the identity matrix I , and $p_{i,j}^n = (P^n)_{i,j}$ refers to the (i, j) entry in P^n . Extending this notation, the $n - step$ transition probability from state i to state j follows:

Definition $\mathbb{P}(X_n = j) = (\lambda P^n)_j$

Definition $\mathbb{P}(X_{n+m} = j | X_m = i) = p_{i,j}^n$

given $(X_n)_{n \geq 0}$ is *Markov*(λ, P) for al $n, m \geq 0$. In other words, the (i, j) entry of $(P^n)_{i,j} = p_{i,j}^n$, gives the probability that the Markov process starting in state i will be in state j after n steps. Constructing a Markov process therefore requires a means of calculating state transition probabilities $p_{i,j}$.

The problem as described thus far, is to define an uncertainty modeling methodology that produces a model capturing the effects of multiple model scales including error of representativeness as well as the regional characteristics of RUC forecasts. In its base form, a Markovian model is capable of representing localized error due to model scales and representativeness based on construction. Imagine a route through

the NAS comprised of a series of waypoints. The distance between these waypoints may vary, allowing for closely grouped waypoints during one phase of the route, and wider spaced waypoints at other points of the route. Research efforts to date, such as those discussed in Section 2.1.2, have considered the correlation intensity of wind speed forecast errors to fit the same exponential function presented in Figure 1. As a result, the correlation of forecast errors is based solely on the distance between two points, not the region in which the waypoints exist; closely grouped waypoints will always exhibit a high correlation of forecast errors regardless of whether or not that characteristic is present. Furthermore, recall that data density is a driving factor in both genesis of model scale errors as well as the correlation of forecast errors along paths through the NAS. The exponential correlation assumption relied on to date must therefore assume equivalent data densities through different regions of the NAS; this idea has been thoroughly contradicted in the RUC literature. Conversely, a Markov model can be constructed such that only sequential forecast error calculations between two specific waypoints are used. One can consider this construction to be analogous to varying the correlation length of the model as a function of route progress and location.

Given a sample route of N waypoints, ACARS flight tracks based on origin-destination airport pair are examined to provide wind speed forecast values sufficiently close to the waypoints comprising the sample route. These forecast error values are then assigned an error state s based on the state space definition S , and state transition probabilities $p_{i,j}$ are calculated by tracking the error states of individual flights as they progress through the sequence of waypoints $\{1, 2, \dots, N\}$ thus forming the state transition matrix P . In order to determine the probable error state at a waypoint n along the sample route, one would simply need to calculate P^n , where n is the n -th waypoint in a sequence of waypoints along the sample route $\{1, 2, \dots, N\}$. This method, however, assumes a homogeneous model. Specifically, the transition

probabilities $p_{i,j}$ are a function of the entire route and not the unique sequence of waypoints comprising the route (i.e.: P is constant). An assumption of homogeneity thus eliminates the need to examine the uncertainty characteristic as a function of both correlation length and region, as the forecast error transition probability characteristics between regions are confounded as are the characteristics of different route leg lengths. As the purpose of the Markovian approach was to examine unique error correlation lengths in different regions of the NAS, an inhomogenous model making use of unique waypoint-to-waypoint state transitions must be constructed.

Given the same series of N waypoints through the NAS and identical forecasts error values clustered by waypoint, a series of unique transition probability matrices $(P^{n,m})_{m \geq n}$ are defined [42]:

Definition $(P^{n,m})_{i,j} = \sum_{\ell \in S} p_{i,\ell}^{n,(m-1)} p_{\ell,j}^m$

where $P^{n,n}$ refers to the identity matrix I . Instead of a single transition probability matrix governing the evolution of the change for all steps $n \in \{1, 2, \dots, N\}$ waypoints, the inhomogeneous chain is constructed using a sequence of transition probability matrices tied to the unique sequence of air waypoints. Consider a sample route of waypoints $\{A, B, C\}$ indexed $\{1, 2, 3\}$, the inhomogenous model consists of three unique transition probability matrices considering forward-only travel: $\{P^{1,2}, P^{2,3}, P^{1,3}\}$. $P^{1,2}$ and $P^{2,3}$ are constructed by examining ACARS flight tracks between waypoints $\{A, B\}$ and $\{B, C\}$ respectively in the same manner as the homogeneous model. $P^{1,3}$ could conceivably be constructed in the same manner by comparing forecast error measurements between waypoints $\{A, C\}$. However, doing so removes the influence of waypoint B on the model. Alternatively, $P^{1,3}$ is determined as the matrix product of the sequence of transitions $P^{1,3} = P^{1,2} P^{2,3} \neq P^{3,2} P^{2,1}$, allowing the chain to evolve as a function of visited waypoints along the route. Forming such transition probability matrices using ACARS flight tracks and archived RUC forecasts is detailed in the following section.

2.3.2 Data Mappings

Background information regarding the RUC forecast product and ACARS meteorological data archives were presented in Sections 2.1.1.1 and 2.1.1.2 respectively. Technical aspects associated with mapping ACARS wind speed measurements into the RUC domain are detailed in this section, along with the methodology used to generate wind speed forecast error measurements that are subsequently used to calculate the Markov transition matrices introduced in Section 2.3.1. The Rapid Update Cycle (RUC) forecast product encompasses a suite of wind speed forecasts varying in horizontal and vertical resolutions. This research effort utilizes the RUC20 forecast, which provides wind speed forecasts on a 20 *km* grid at 50 vertical levels (a full specification is provided by in [5]).

2.3.2.1 Horizontal Resolution

The RUC20 provides data in the form of gridded arrays across 50 vertical levels. These arrays map to a specific spatial domain defined by a Lambert conformal projection, referred to as the Advanced Weather Interactive Processing System (AWIPS) 252 grid. A Lambert conformal projection seats a cone over a geodetic datum spheroid representing the Earth and projects from the Earth datum to the cone conformally, meaning angles are preserved. The cone is then “unrolled” resulting in the projected surface. The details of the projection method itself are not presented, as it is beyond the scope of this work. However, there is sufficient literature on the topic available to the reader. There are also several industry standard open source computational libraries such as PROJ.4 available for this task. The World Geodetic System (WGS) 84 Earth specification is used as the geodetic datum for the Lambert projection; WGS84 is the same datum used by the Global Positioning System. The AWIPS 252 grid provides total coverage of the continental United States across its 301×225 domain, with a nominal grid length of 20.317 *km* at 25° N latitude; the extents of the

AWIPS 252 grid are illustrated in Figure 2 [64] and the parameters of the Lambert conformal projection are provided in Table 1. The reader should note the differences between the East and North directions measured relative to the parallels meridians as compared to the grid-relative East and North directions, referred to herein as the \mathbf{u} and \mathbf{v} directions. At any point in the grid, if one were to travel due East along a parallel, this travel would correspond to a projected path of travel in both the \mathbf{u} and \mathbf{v} directions, whereas the Earth-centric travel would feature no Northern component. This effect is a direct result of the varying map-scale factor of the projection and becomes more significant as one approaches the projection boundaries. Consequently, any vector values measured in an AWIPS 252 centric coordinate system must first be rotated before they are compared to East-North centric vector values. Wind speed values in a RUC20 grid are provided at grid vertices in terms of \mathbf{u} and \mathbf{v} unit directional components. Contrary to meteorological convention in regards to wind directional notation, the resultant vector giving wind speed and direction:

$$\mathbf{w}_{forecast} = u \cdot \mathbf{u} + v \cdot \mathbf{v}$$


points *in the direction of the wind*, as opposed to the conventional notation in which the directional vector points in the direction from which the wind is blowing. Finding the forecast wind at a local point (lat, lon) thus requires interpolating the wind speed values at the four grid vertices surrounding the local point. First, the point (lat, lon) is mapped into the AWIPS 252 domain using the Lambert conformal projection defined in Table 1 producing a point (x, y) . The neighboring grid vertices and corresponding wind speed forecast values are easily found:

$$n_1 = (\lfloor x \rfloor, \lfloor y \rfloor) \rightarrow \mathbf{w}_1 = u_1 \cdot \mathbf{u} + v_1 \cdot \mathbf{v}$$

$$n_2 = (\lfloor x \rfloor, \lceil y \rceil) \rightarrow \mathbf{w}_2 = u_2 \cdot \mathbf{u} + v_2 \cdot \mathbf{v}$$

$$n_3 = (\lceil x \rceil, \lceil y \rceil) \rightarrow \mathbf{w}_3 = u_3 \cdot \mathbf{u} + v_3 \cdot \mathbf{v}$$

$$n_4 = (\lceil x \rceil, \lfloor y \rfloor) \rightarrow \mathbf{w}_4 = u_4 \cdot \mathbf{u} + v_4 \cdot \mathbf{v}$$



NCEP Grid 252

Figure 2: The AWIPS 252 grid.

Table 1: AWIPS 252 grid projection parameters.

Parameter	Value
N_x	301
N_y	225
D_x	20.317625 km
D_y	20.317625 km
Longitudinal Axis	-95° E
Latin 1	25° N
Latin 2	25° N
Grid Corner (1,1)	(16.281° N, 126.138° W)
Grid Corner (1,225)	(54.127° N, 139.856° W)
Grid Corner (301,225)	(55.481° N, 57.383° W)
Grid Corner (301,1)	(17.340° N, 69.039° W)

A bilinear interpolation is then used to find the \mathbf{u} and \mathbf{v} components of the wind at the local point (x, y) :

$$\begin{aligned}\mathbf{u}_{x,y} &= [u_1(1 - \Delta x)(1 - \Delta y) + u_2(1 - \Delta x)\Delta y + u_3\Delta x\Delta y + u_4\Delta x(1 - \Delta y)] \cdot \mathbf{u} \\ \mathbf{v}_{x,y} &= [v_1(1 - \Delta x)(1 - \Delta y) + v_2(1 - \Delta x)\Delta y + v_3\Delta x\Delta y + v_4\Delta x(1 - \Delta y)] \cdot \mathbf{v}\end{aligned}$$

where $\Delta x = x - \lfloor x \rfloor$ and $\Delta y = y - \lfloor y \rfloor$. Due to the varying map scale of the Lambert conformal projection and its effects on \mathbf{u} and \mathbf{v} relative to East and North, it is necessary to provide a further correction with respect to the wind speeds at the neighboring vertices. First, the wind speed at the local point (x, y) is interpolated using the wind speeds at the neighboring vertices:

$$s_{x,y} = |\mathbf{w}_1|(1 - \Delta x)(1 - \Delta y) + |\mathbf{w}_2|(1 - \Delta x)\Delta y + |\mathbf{w}_3|\Delta x\Delta y + |\mathbf{w}_4|\Delta x(1 - \Delta y)$$

A correction factor is then calculated:

$$K_s = \frac{s_{x,y}}{|\mathbf{u}_{x,y} + \mathbf{v}_{x,y}|}$$

The local wind speed vector at a point (x, y) follows:

$$\mathbf{w}_{x,y} = K_s(\mathbf{u}_{x,y} + \mathbf{v}_{x,y}) \quad (6)$$

The result derived in Equation 6 returns the local wind speed at a point (x, y) in terms of the unit directional vectors \mathbf{u} and \mathbf{v} relative to the Lambert conformal projection of the AWIPS 252 grid. A further transformation is required in order to represent the local wind speed at (x, y) in terms of East and North directional components \mathbf{e} and \mathbf{n} relative to the original candidate point (lat, lon) [25]. The rotational constant is first defined in terms of the AWIPS 252 projection parameters:

$$r = \sin(25^\circ)$$

where 25° N is the Latin 1 projection tangent at which the AWIPS 252 projection is true. The local rotation angle is then calculated based on the distance from the

local point (lat, lon) from the longitudinal axis of the AWIPS 252 grid (the meridian aligned with the Cartesian x-axis):

$$\theta_2 = r(lon - (-95))$$

The local wind speed at a point (lat, lon) in terms of local East and North directional components \mathbf{e} and \mathbf{n} follows:

$$\begin{aligned} e_{lat,lon} \cdot \mathbf{e} &= \cos(\theta_2)(\mathbf{w}_{x,y} \cdot \mathbf{u}) + \sin(\theta_2)(\mathbf{w}_{x,y} \cdot \mathbf{v}) \\ n_{lat,lon} \cdot \mathbf{n} &= -\sin(\theta_2)(\mathbf{w}_{x,y} \cdot \mathbf{u}) + \cos(\theta_2)(\mathbf{w}_{x,y} \cdot \mathbf{v}) \end{aligned}$$

leaving the local wind speed at a point (lat, lon) :

$$\mathbf{w}_{lat,lon} = e_{lat,lon} \cdot \mathbf{e} + n_{lat,lon} \cdot \mathbf{n} \tag{7}$$

The reader should again note that the vector result presented in Equation 7 provides a vector result that points *in the direction of the wind*.

2.3.2.2 Vertical Resolution

RUC20 gridded data is available at 50 vertical levels organized by pressure layer relative to a standard atmosphere. Calculating the wind speed at a specific location $(lat, lon, altitude)$ is simply a matter of mapping the desired pressure altitude (or flight level) to the corresponding pressure in a standard atmosphere, locating the two neighboring data grids surrounding the local pressure altitude, interpolating the local wind speeds longitudinally in each of the neighboring data grids using the method defined in section 2.3.2.1, then linearly interpolating on pressure altitude between the layers.

2.3.2.3 Direction and Time

To this point, two facets of the forecast error calculation have not been discussed: the direction vs. magnitude of forecasted and actual wind speeds, and the time reference of the RUC20 forecast releases. Wind speeds are expressed as vector quantities

with a speed and direction. A Markov model describing the uncertainty in a wind speed forecast model must necessarily be bivariate in order to encompass both values. Accordingly, this research effort will generate two individual models in terms of the Eastern component of the wind speed forecast error along the \mathbf{e} unit direction defined in Section 2.3.2.1, and the corresponding Northern component of the forecast error along the \mathbf{n} unit direction also defined in Section 2.3.2.1. As a result, the Eastern and Northern components of the Markov uncertainty model are inherently decoupled. As the end use of the model is longer term trajectory prediction in the en route regime, and local vorticity (the mixing of directional components of the wind) is largely a localized effect, a decoupled model should provide reasonable results as any ignored vorticity will not compound over the course of a tracked flight route.

The time component of RUC20 forecast releases also provides another interesting consideration. Forecast release time and relative forecast age were introduced in Section 2.1.2 while discussing ensembled forecasts. As a review, an individual forecast has a release time and a valid time relative to the forecast release. For example, a RUC20 forecast released at 00:00 GMT will contain wind speed forecasts in one hour increments until 06:00 GMT, leaving six valid forecast times for the 00:00 GMT forecast release (i.e.: the 00:00 GMT 1-hour forecast is valid from 00:00 GMT to 1:00 GMT, the 00:00 GMT 2-hour forecast is valid from 1:00 GMT to 2:00 GMT, etc.). The rolling structure of the RUC20 forecast cycle two extreme paths through the time history of RUC20 releases. The first path holds the forecast release static, considering each of the hourly forecasts associated with that release to provide $\mathbf{w}_{forecast}$ values for the model. This scenario can be considered analogous to a pilot loading a wind speed forecast before taking off and not updating the wind speed forecasts at any point along the route. The second path only considers the 1-hour release of each new RUC20 release. Every hour, the freshest forecast release is used to provide a one hour forecast. Conversely, this scenario is analogous to a pilot updating the forecast wind

speeds as soon as the newest RUC20 forecast is released. From this point forward, a model using the former scenario will be referred to as the “base” case, and a model using the latter scenario will be referred to as the “updated” case.

2.3.3 A Sample Construction

As discussed at the beginning of Section 2.3, air carriers tend to follow common routes through the NAS between origin and destination airports. This practice results in large concentrations of ACARS flight tracks (and ultimately forecast error measurements) along similar ground tracks between origin and destination airport pairs. Given the methodologies defined in Sections 2.3.1 and 2.1.1, it is possible to construct an inhomogeneous Markovian RUC forecast uncertainty model along a common route, such as the model for the route between Seattle’s Tacoma International Airport (KSEA) and Hartsfield-Jackson International Airport in Atlanta (KATL) that is presented in this section. First, the route itself as well as data considerations are discussed; calculation of the transition probability waypoints between the route waypoints then follows; a visualization of the Markov chain probabilities and a simulated wind scenario along the route are then discussed in closing.

2.3.3.1 The Data Set

In order to generate a sufficient number of usable data points to calculate Markov transition probabilities, a large volume of RUC-matched data pairs is required. Accordingly, this study considers three years of ACARS flight tracks between each origin and destination pair, recorded between January 1, 2008 and December 31, 2010. However, the data set must first be sanitized to ensure only data points meeting the following criteria are included:

1. Measurement occurs during the cruise portion of the flight.
2. Measurement passes all ACARS quality control checks for temporal and spatial

validity.

3. Measurement is recorded while the aircraft maintains a roll angle of less than 5° as advised by Biseaux et al. [11]

Sanitation of the data set leaves a total of 2,797 remaining flights tracks for the sample scenario of a flight from KSEA to KATL for the three year period. A further consideration must be lent, however, to the pairwise proximity of the remaining flight tracks. Though commercial operators tend to follow similar routings between an origin and destination pair, exogenous factors such as storms or congestion may lead to deviation from the standard track. The RUC-matched ACARS data density as a function of RUC grid cell is illustrated in Figure 3, where the heat of the individual cell is a function of how many ACARS data points fall within each individual cell. As a band of flight tracks exists along the Northern most edge of the KSEA to KATL flight tracks, the sample route is constructed such that the route follows a ground track through this data-dense region. Constructions for routes between other origin and destination airport pairs must follow a similar methodology; the data available along routes between the pair is first examined, then a candidate path is constructed through data dense regions. Though not a universal solution, this approach both ensures sufficient data coverage for calculating the Markov chain as well as satisfies the intended end-use case of the uncertainty model for an aircraft in cruise.

ACARS Record Density - KSEA to KATL Route

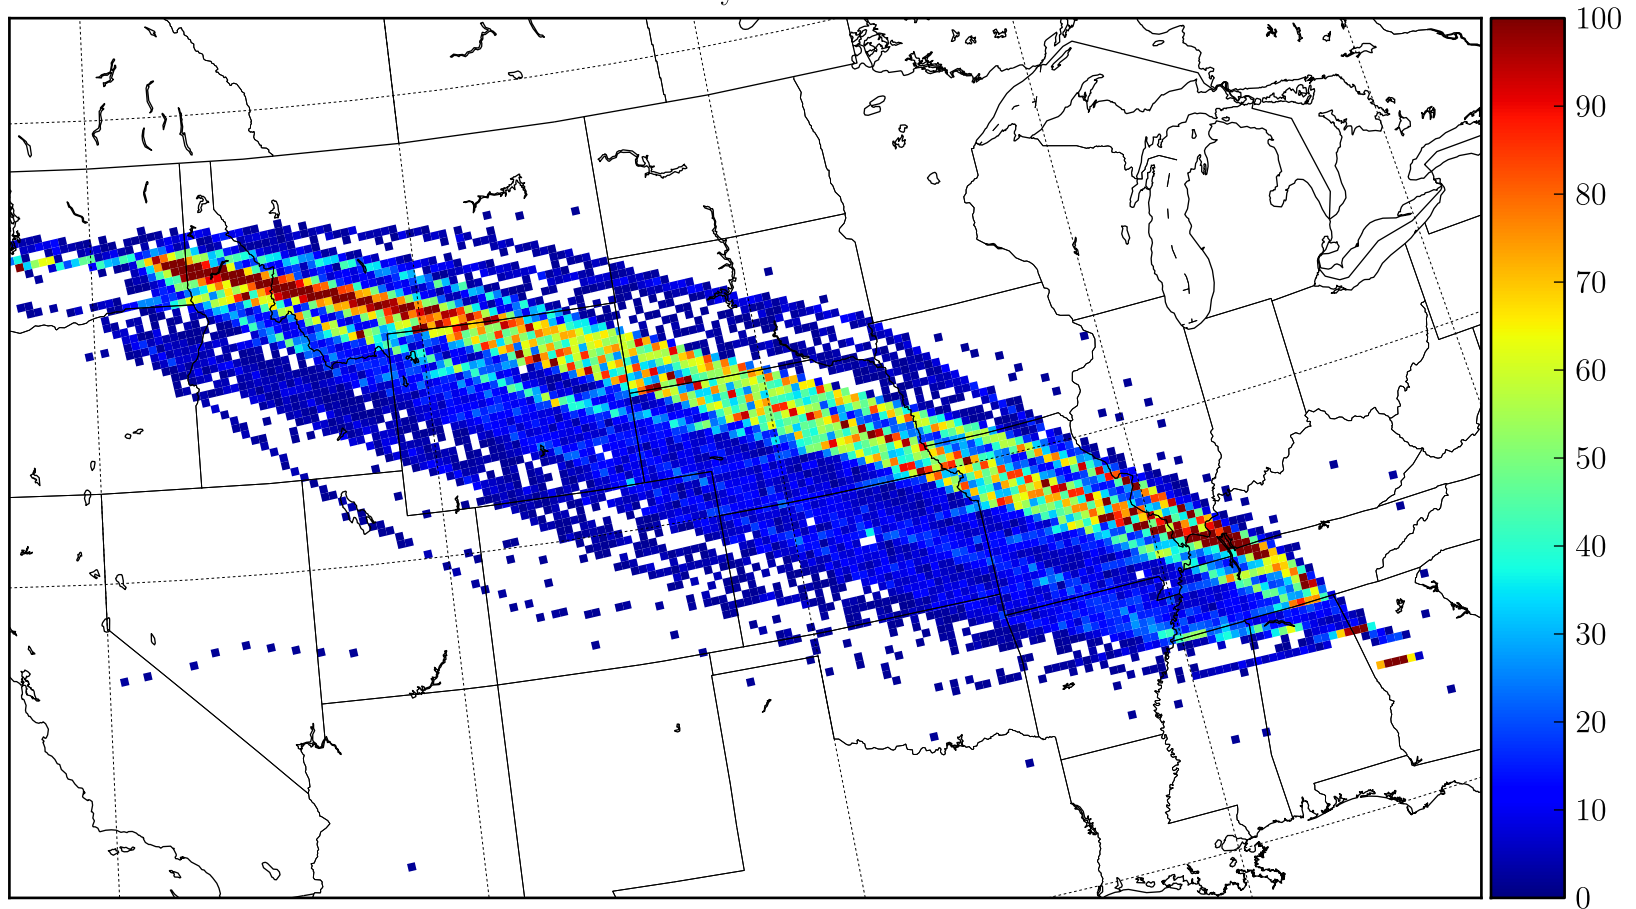


Figure 3: RUC-matched ACARS data density in RUC cells along the KSEA to KATL route.

Table 2: Sample route used to construct uncertainty model for flights between KSEA and KATL.

Waypoint	Latitude ($^{\circ}N$)	Longitude ($^{\circ}E$)	Length (nm)	Heading ($^{\circ}$)
MWH	47.2108633	-119.3168167	306.790	102.7
HIA	45.8616667	-112.1697222	224.602	104.0
SHR	44.8422222	-107.0611111	71.503	114.0
GCC	44.3477778	-105.5436111	265.048	111.8
ANW	42.5691667	-99.9897222	176.084	123.0
LNK	40.9238889	-96.7419444	101.298	124.2
STJ	39.9605556	-94.9252778	259.156	120.4
FAM	37.6733333	-90.2341667	194.014	117.2
BNA	36.1369722	-86.6847778	N/A	N/A

2.3.3.2 Lateral Path Construction

Though not necessarily a function of the forecast uncertainty model, route construction must first be discussed in order to give context to the end-use case of the model. As the model is position-based and tied to a specific ground track, considerations regarding the construction of the track itself should first be explored as they will ultimately impact the model itself. Flight routes are defined in an aircraft’s flight plan as a series of latitude and longitude tuples describing the location of standardized waypoints or navigation aids. The flight plan does not describe the intended lateral path of the aircraft exactly, but rather serves as a series of locations the aircraft must visit (either closely or exactly) along its route. The flight management system calculates the aircraft’s intended ground track given the route waypoints supplied by the flight plan, effectively constructing the entire route to be flown by the aircraft on a point-to-point basis. A route for the sample scenario of a flight between KSEA and KATL is presented in Table 2 as a series of waypoints, and is further illustrated in Figure 4.

Sample Route - KSEA to KATL

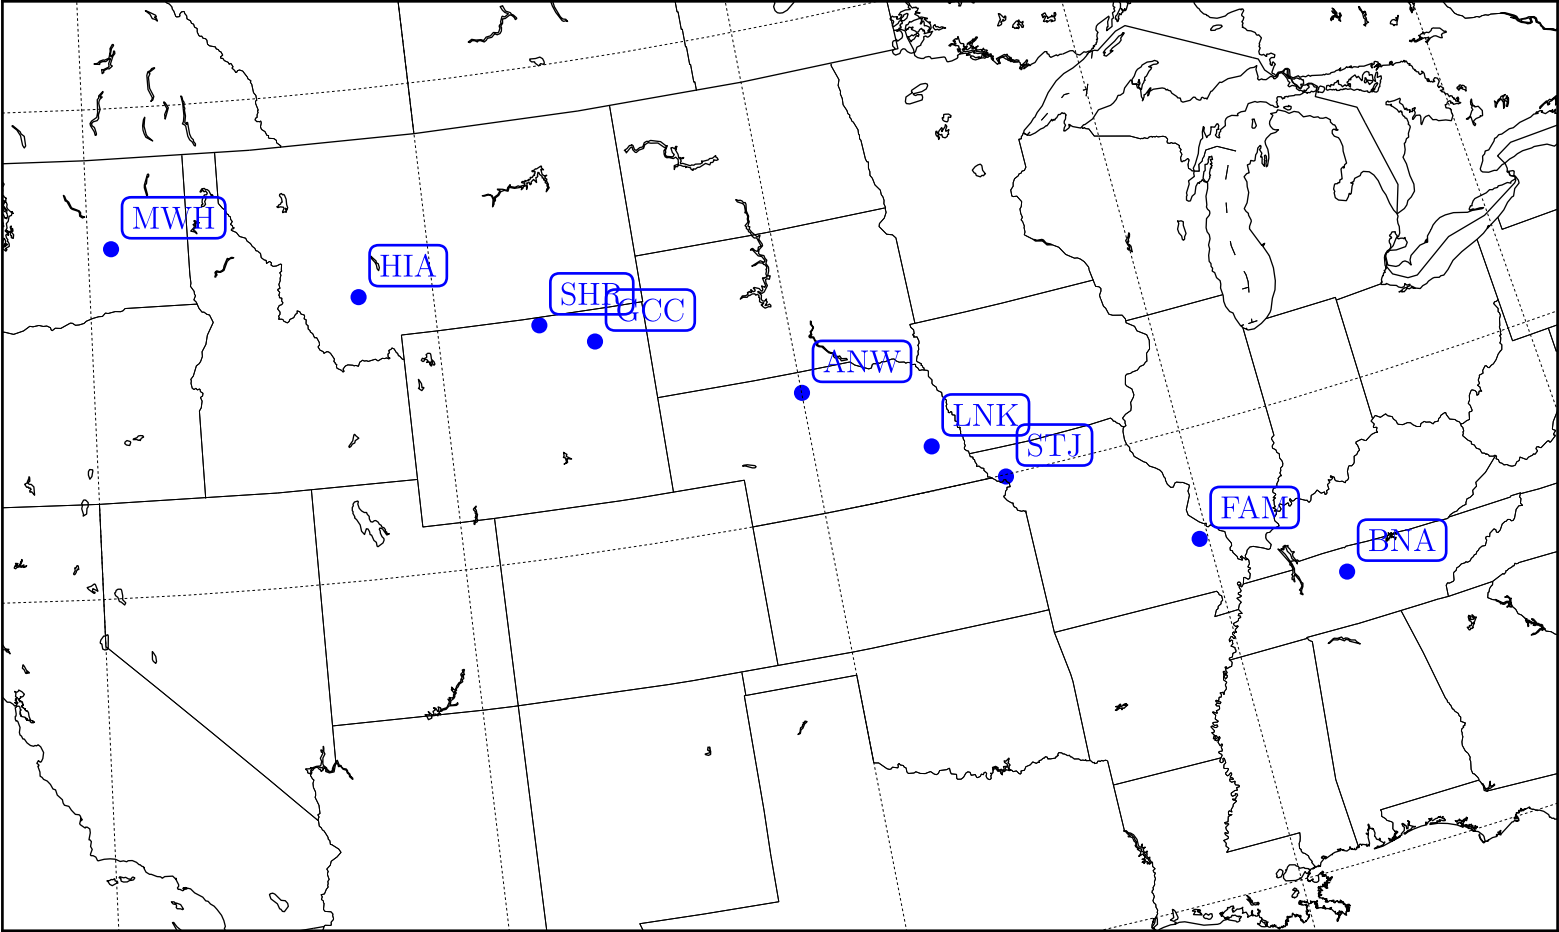


Figure 4: KSEA to KATL sample route in terms of standard waypoints.

Table 3: Sample route segments used to construct uncertainty model for flights between KSEA and KATL.

Start	Finish	Shortest Length (nm)	Heading ($^{\circ}$)
MWH	HIA	306.790	102.7
HIA	SHR	224.602	104.0
SHR	GCC	71.503	114.0
GCC	ANW	265.048	111.8
ANW	LNK	176.084	123.0
LNK	STJ	101.298	124.2
STJ	FAM	259.156	120.4
FAM	BNA	194.014	117.2

Though not explicitly required, it is often desirable for an aircraft to follow the shortest path between consecutive waypoints, roughly a great circle path. Recall that a great circle divides a sphere into two equal hemispheres. A great circle path between two points on the surface of a sphere is the minor arc connecting the two points along the diameter of a great circle; in reference to spherical geometry, a great circle path is a straight line between two points lying on the surface. The Earth, however, cannot be precisely represented in terms of a sphere, but rather as an oblate spheroid per the WGS84 datum. Vincenty’s algorithm provides an iterative method for calculating analogous great circle paths between two points on the surface of an oblate spheroid [87]. Table 3 lists the great circle distances and headings along individual segments of the sample route previously defined in Table 2 and Figure 4.

The sample scenario flight path is now defined in terms of a series of great circle segments between defined waypoints through the data dense regions depicted in Figure 3. Again, it should be mentioned that an aircraft isn’t required to maintain a great circle path between waypoints given traffic, weather, or other outside considerations. Though, given the assumption of a fixed flight path with a predefined flight plan, assuming a great circle arc between waypoints is reasonable. Further consideration must be given to heading changes along the route, however. As mentioned previously, the flight management system constructs the aircraft’s exact ground track




Figure 5: FAA illustration of fly-by and fly-over waypoints.

given the series of waypoints defined in the flight plan. To this point, that path has been defined as a series of great circle arcs connecting consecutive waypoints. Any heading change along the path is handled instantaneously meaning the aircraft flies an infeasible sharp corner. In practice, aircraft flight management systems handle heading changes in one of two manners based on the turn waypoint. Turns at a “fly-by” waypoint are initiated prior to the aircraft reaching the waypoint allowing the aircraft to perform a steady turn effectively ”cutting the corner” at the waypoint. Conversely, “fly-over” waypoints are first flown over before a series of turns is executed to bring the aircraft back on course. Both scenarios are illustrated in Figure 5 [30]. In the context of Figure 5, an instantaneous heading change implies that the aircraft follows the nominal track exactly and does not shorten or lengthen the path when executing a turn. In high fidelity trajectory prediction applications, turn modeling methodology heavily impacts predicted flight times. Large heading changes may significantly stretch or shorten paths depending on both the type of waypoint and the aircraft’s cruise speed leading to inaccurate flight time estimation if not accounted for. Mondoloni’s study of turn modeling errors demonstrated that along-track distance (the length of the path to be flown) error rose as high as 2.71 *nm* for a 90° heading change at 400 *kts* cruise speed [63]. The scenarios considered in this study, however, do not present such stark heading changes as the scenario posed

by Mondoloni or the scenario depicted in Figure 5. On the contrary, in the case of the sample scenario of a flight between KSEA and KATL, the heading difference between neighboring flight segments is on the order of a few degrees, with the largest heading change of 11.2° occurring at ANW. For small heading changes on the order of 20° , the same Mondoloni study found along track error on the order of 0.02 nm . As a result, this research effort will only consider routes with negligible heading changes in order to reasonably approximate turns as occurring instantaneously.

A final consideration must be lent to the segment distances listed in Table 3. The segment lengths for the sample scenario range from 71.503 nm to 306.790 nm . There is nothing inherently wrong with this construction from a flight path standpoint as it is simply the most convenient means of defining a specific path through the desired region. However, as the methodology discussed Section 2.3.1 defines Markov transition probabilities on transitions from waypoint to waypoint, the Markov chain will evolve over significantly different distances as a function of route progress. In other words, adhering to the predefined route waypoints from a typical aircraft flight plan produces a model with irregular along-track resolution. It is therefore desirable to recast the flight path in terms of intermediate or artificial waypoints at equidistant points along the original path defined in Tables 2 and 3. The sample route in terms of the original waypoints as well as 15 artificial waypoints is depicted in Figure 6, creating flight segments roughly 100 nm long. The new route reads as a blended string of waypoints new and old $\{\text{MWH, A1, A2, A3, HIA, A4, A5, SHR, } \dots \}$; the full list of flight segments with artificial waypoints can be found in Table 4. The path defined in Figure 6 and Table 4 follows the original sample route defined in Table 2, Table 3, and Figure 4 exactly, with the artificial waypoints simply representing 100 nm distance markers along the original route. 100 nm was chosen as a reasonable discretization distance based on industry guidance regarding the capabilities of modern flight management systems [Personal Communication. Dr. Liling Ren, GE

Table 4: Sample route segments with artificial waypoints for KSEA to KATL scenario.

Start	Finish	Shortest Length (nm)	Heading ($^{\circ}$)
MWH	A1	99.906	102.7
A1	A2	99.906	104.4
A2	A3	99.906	106.1
A3	HIA	7.0716	107.7
HIA	A4	92.834	104.0
A4	A5	99.906	105.5
A5	SHR	31.862	107.1
SHR	A6	68.044	114.0
A6	GCC	3.459	115.0
GCC	A7	96.447	111.8
A7	A8	99.906	113.2
A8	ANW	68.695	114.7
ANW	A9	31.211	123.0
A9	A10	99.906	123.4
A10	LNK	44.966	124.6
LNK	A11	54.938	124.2
A11	STJ	46.340	124.8
STJ	A12	53.546	120.4
A12	A13	99.906	121.1
A13	A14	99.906	122.2
A14	FAM	5.798	123.3
FAM	A15	94.108	117.2
A15	BNA	99.906	118.3

Global Research Center, May 2011]. For example, waypoints A7 and A8 lie on the great circle path between GCC and ANW. A9 and A10 on the other hand lie on the great circle path between ANW and LNK. The path between A8 and A9 through ANW is not a great circle path due to the heading change at ANW as required by the original flight plan defined in Table 2. However, the total length of the great circle paths from A8 to ANW and ANW to A9 sum to the desired path length. Using the path waypoints defined in Table 4, Markov transition probabilities can be calculated at equidistant track lengths along the originally planned flight path. It should be apparent to the reader that the Markov transitions will be calculated at the artificial waypoints $\{A1, A2, \dots, A15\}$ rather than the original planned waypoints in order to

preserve along track distance and in turn, along-track model resolution.

Artificial Route - KSEA to KATL

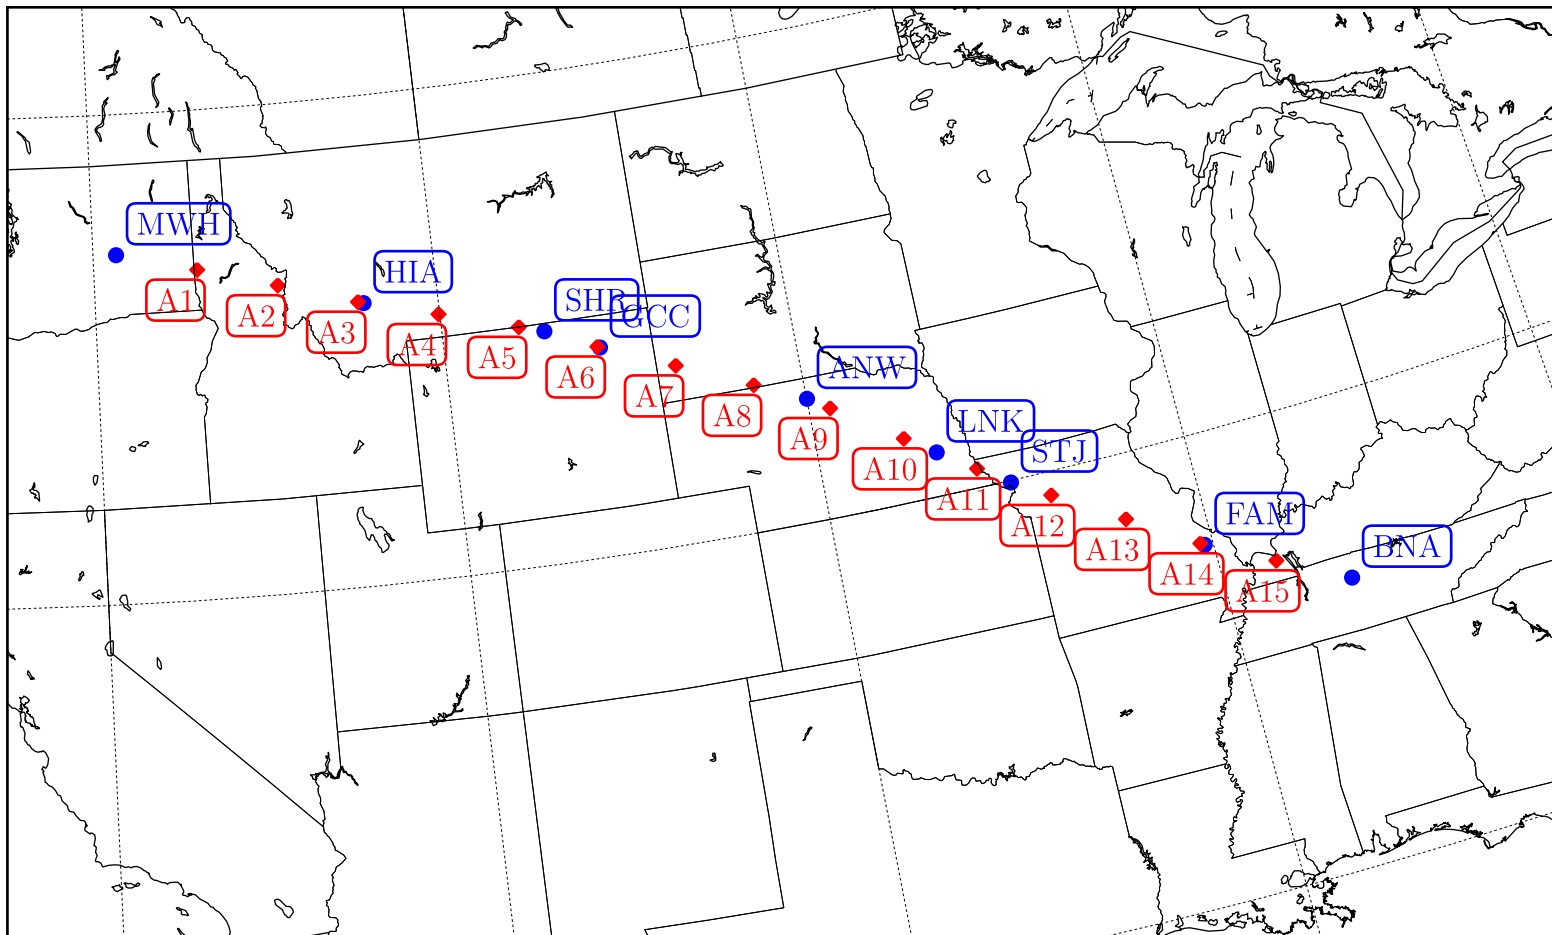


Figure 6: KSEA to KATL sample route in terms of artificial waypoints.

2.3.3.3 *Altitude and Speed*

Section 2.3.3.2 defined the aircraft’s lateral flight path in terms of a series of great flight segments between air waypoints defined in the aircraft’s flight plan. To this point, altitude or cruise speed have not been mentioned. Along with lateral path waypoints, a flight plan includes desired flight altitudes and speeds at each waypoint or cruise-climb segments where available. This specification provides an additional factor for consideration generating a Markov model: the altitude layer of the model itself.

The currently formulated model includes aircraft flight data from all available altitude layers, forming a single model regardless of cruise altitude. The set used to generate Figure 3, for example, includes data points ranging in altitudes from 35,000 *ft* to 41,000 *ft*. Upon first inspection, it may seem necessary to further subdivide the 2,797 sanitized flight tracks into altitude layers in terms of flight level, where flight level is the barometric pressure altitude of the standard atmosphere in terms of hundreds of feet (i.e.: 35,000 *ft* pressure altitude corresponds to flight level 350). However, prior work by the author determined that wind forecast error values were more sensitive to direction of travel rather than altitude alone [86]. First recall that Federal Aviation Regulation (FAR) Part 91 defines cruise altitudes for aircraft operating above flight level 290 based on direction of travel [31]:

- “(i) On a magnetic course of zero degrees through 179 degrees, any odd flight level, at 2,000-foot intervals beginning at and including flight level 290 (such as flight level 290, 310, 330, 350, 370, 390, 410); or
- (ii) On a magnetic course of 180 degrees through 359 degrees, any even flight level, at 2000-foot intervals beginning at and including flight level 300 (such as 300, 320, 340, 360, 380, 400).”

Part 91 requires that all East-bound air traffic maintain an “odd” flight level starting at flight level 290, and all West-bound traffic maintains an “even” flight level starting at flight level 300, meaning that air traffic alternates in 1000 *ft* intervals based on direction of travel. The author’s 2009 study demonstrated that the RUC20 forecast error values were nearly statistically identical within directional subgroups. East-bound flight tracks odd flight levels exhibited nearly identical error properties regardless of flight level; this result held across the data for West-bound flight levels. Given the results of this prior study, subdividing the model further in terms of flight level is unwise, as the added fidelity of a per flight level model is negligible compared to the severe drop in available flight tracks when the data set subdivided into altitude layers.

2.3.3.4 Error Calculation

A sample route between KSEA and KATL was defined through a data dense region of the airspace along common routings between the two airports in Section 2.3.3.2. Artificial waypoints were then calculated along the route to provide consistent horizontal resolution along the sample route’s lateral path. In order to generate Markov transition probabilities between the artificial waypoints, forecast error values at each waypoint must be calculated by comparing the RUC wind speed forecasts at the waypoints with archived ACARS wind speeds recorded sufficiently close to the waypoints.

The closeness of ACARS wind speed measurements to a specific waypoint provides an interesting challenge. Not only do individual measurements need to be within some neighborhood of the waypoint, but this condition must hold for all waypoints along the string of measurements for a given ACARS track. In essence, the entire ACARS ground track must be sufficiently close to the desired flight route in order to be used. A technique considering the nearest approaches of ACARS flight routes to the desired flight path was developed to handle such filtering:

Algorithm 1 Filters ACARS ground tracks close to a desired route.

```
for all ACARS tracks in set  $\{1 \dots j \dots M\}$  do  
  for all Waypoints along desired route  $\{1 \dots i \dots N\}$  do  
     $d_{j,i} \leftarrow$  distance of nearest approach between ACARS track  $j$  and waypoint  $i$   
  end for  
   $\bar{d}_j \leftarrow \frac{1}{N} \sum_{i=1}^N d_{j,i}$  {mean distance of nearest approach}  
  if  $\bar{d}_j \leq 40km$  then  
    Accept ACARS ground track  $j$ .  
  else  
    Reject ACARS ground track  $j$ .  
  end if  
end for
```

Algorithm 1 uses the mean distance of closest approach between individual ACARS flight tracks in a set and a series of waypoints specifying a desired route. If the mean distance of closest approach for a single ACARS track is less than 40 *km*, the ACARS track is considered sufficiently close to the desired flight route. Essentially, the algorithm searches for routes that are on average within one RUC cell of the desired route. Algorithm 1 was used to filter the original 2,797 flights between KSEA and KATL with respect to the sample route constructed previously; the resulting flight track are presented in Figure 7.

Artificial Route with Filtered ACARS Tracks - KSEA to KATL

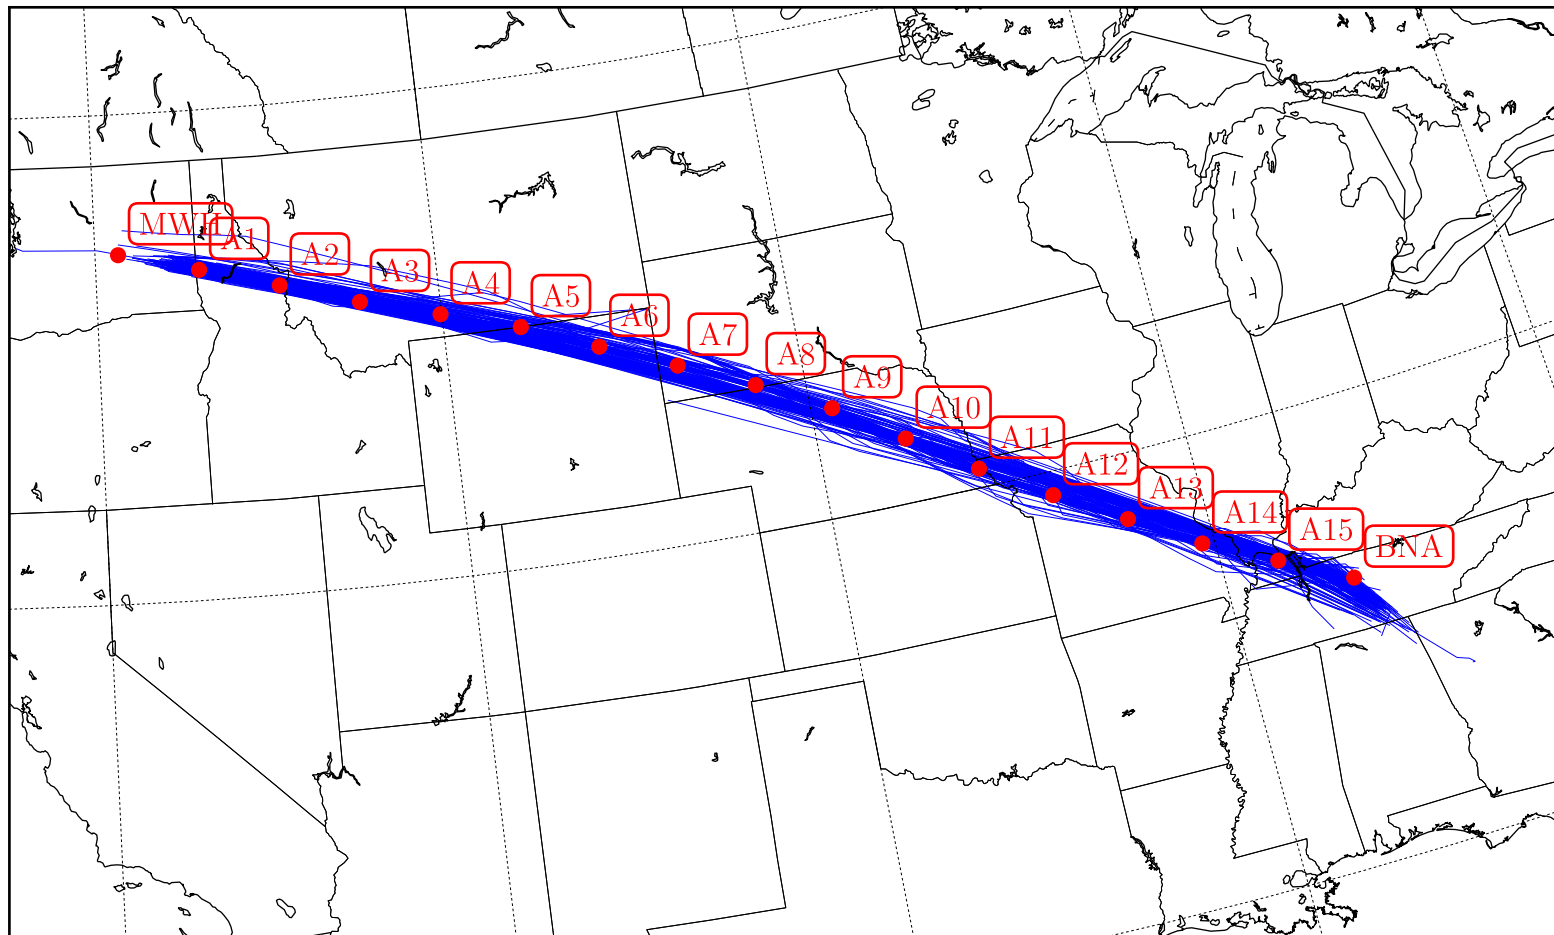


Figure 7: Filtered ACARS tracks along KSEA to KATL sample route.

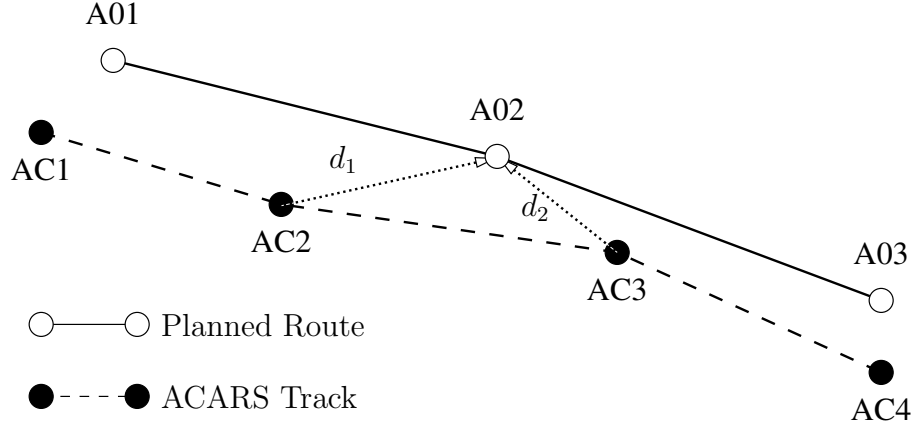


Figure 8: Wind Interpolation scenario.

Though steps have been taken to minimize deviation between the archived ACARS tracks and the proposed sample track by including only similar ground paths in the study, further consideration must be lent to the location of the ACARS measurements themselves. Ideally, the ACARS archives would provide wind speed measurements from flights passing over the same waypoints as the planned sample route. However, as previously demonstrated this is not the case. Consider the case of a planned route consisting of a series of waypoints $\{A1, A2, A3\}$ alongside a neighboring ACARS flight track consisting of wind speed measurements at locations $\{AC1, AC2, AC3, AC4\}$ as presented in Figure 8. A means of extrapolating the wind speed at waypoint A02 is needed given the available ACARS measurements at known locations $\{AC1, AC2, AC3, AC4\}$. [50] provides a survey of meteorological extrapolation and interpolation routines given a variety of data structures, a few of the suggested routines being function fitting methods such as cubic spline interpolation [70] and Kriging [48], successive corrective methods, and Kalman filtering. However, the use case defined to this point (and illustrated in Figure 8) requires that wind speed forecast error be tracked on a flight-by-flight basis, meaning that only the contributions from a single string of ACARS measurements may be used to generate a single wind speed error measurement at waypoint A02. [50] suggests that inverse distance weighting [81] be used in the case of such sparse data sets. Inverse distance weighting simply calculates

a weighted sum of known observations, where the weights are specified based on the distance of the known observations from the desired location. Generally, the value of a field f_a at a point i can be approximated using a combination of K observations f_o weighted by a function of the inverse of the distance from the point i to the k -th observation:

$$f_a(\mathbf{r}_i) = \sum_{k=0}^K \frac{w(\mathbf{r}_{i,k})f_o(\mathbf{r}_k)}{\sum_{j=0}^K w(\mathbf{r}_{i,j})} \quad (8)$$

$$w(\mathbf{r}_{i,k}) = \frac{1}{d(\mathbf{x}_i, \mathbf{x}_k)^p} \quad (9)$$

where $d(\mathbf{x}_i, \mathbf{x}_k)$ is the distance operator which returns the distance between points i and k , and p is the power parameter (for sparse data sets, [50] suggests using $p = 2$). It should be noted that distance weighting schemes are not appropriate for scenarios involving larger data sets, as they are heavily directionally dependent or anisotropic. Inverse distance weighting also fails to capture the higher order effects quantified by more robust schemes. However, given the sparseness of the available data set, there is no advantage to pursuing such methods.

One can use the method posed in Equations 8 and 9 to approximate the ACARS-measured wind speed at waypoint A02 as a combination of the recorded measurements at locations AC2 and AC3:

$$\mathbf{w}_{A02} = \frac{\left(\frac{1}{d_1^2}\right) \mathbf{w}_{AC2}}{\frac{1}{d_1^2} + \frac{1}{d_2^2}} + \frac{\left(\frac{1}{d_2^2}\right) \mathbf{w}_{AC3}}{\frac{1}{d_1^2} + \frac{1}{d_2^2}} \quad (10)$$

where d_1 and d_2 are the distances between locations AC2 and waypoint A02 and AC3 and waypoint A02 respectively, and \mathbf{w}_{AC2} and \mathbf{w}_{AC3} are the vector wind speeds at locations AC2 and AC3 respectively. It was found that inverse weighting is very sensitive to the distances between observations given the suggested power factor $p = 2$. As a result, only the two closest ACARS observation to a specified waypoint are used to interpolate the ACARS wind speed at the waypoint; adding more points to the interpolation provided little benefit given the distances between ACARS observations.

The reader may have noted that the expression derived in Equation 10 is in terms of vector quantities \mathbf{w}_{A02} , \mathbf{w}_{AC2} , and \mathbf{w}_{AC3} . As previously mentioned in Section 2.3.2, ACARS wind speed observations are reported in terms of a wind speed and wind direction, where the direction value is reported in terms of a 0° true-North reference and signifies the direction *from which the wind is blowing*. In order to ensure ACARS wind speed observations are expressed in the same frame of reference as the RUC wind forecasts, ACARS winds speed observations are further transformed to a component-based form:

$$\mathbf{w}_{lat,lon} = e_{lat,lon} \cdot \mathbf{e} + n_{lat,lon} \cdot \mathbf{n}$$

in terms of local East and local North vector components at a point (lat, lon) , where \mathbf{e} and \mathbf{n} point *with the direction of the wind* (effectively reversing the direction of the ACARS wind direction observation). After transforming ACARS observations into a common frame of reference as the RUC wind speed forecasts, inverse distance weighting (Equations 8 and 9) is used to generate ACARS wind speed observations in terms of East and North components at each of the artificial waypoints along the sample route $\{MWH, A1, \dots, BNA\}$. Using the methods outlined in Section 2.3.2, the ACARS observations are then compared to the corresponding RUC forecasts given the *forecast – actual* convention defined in Equation 5 creating a series of wind speed forecast error measurements as a function of individual flights at each of the sample route waypoints.

2.3.3.5 Error Distributions

RUC forecast errors have been calculated at each of the artificial waypoints constructed along the sample route between KSEA and KATL by comparing ACARS wind speed observations from flights along similar routes with the corresponding RUC forecasts on a per flight basis. Referring to the model definition constructed in Section 2.3.1, a series of $q = 11$ error states $s \in S$ are now defined on the range $R = [-11, 11]$,

where S contains the error state intervals $\{(-11, -9], (-9, -7], \dots, (9, 11]\}$. This range and number of states were chosen to provide enough granularity while ensuring adequate data density and a zero-centered model. The RUC forecast errors calculated at each waypoint are now mapped into the state space S , and a normalized histogram of the error populations generated. Figures 9 and 10 illustrate two such histograms for the RUC forecast errors at MWH and A1. In each histogram, the states $s \in S$ are listed along the horizontal axis, with the vertical axis representing the percentage of the total error values mapped to a particular state s_i . Error counts in both the local East \mathbf{e} and local North \mathbf{n} directions are plotted alongside one another. The reader should further note that figures 9 and 10 also represent forecast error counts for the base model case discussed in Section 2.3.2.3.

As standalone results, the forecast error distributions presented in Figures 9 and 10 provide interesting results. The distributions in both directions are roughly zero mean, and there is a slight tendency at both waypoints for there to be a higher percentage of peak errors in the positive direction rather than negative, meaning the RUC forecast is slightly more likely to demonstrate higher over predictions rather than under predictions based on the *forecast - actual* convention. This result is further illustrated in Figures 11 and 11 depicting cumulative histograms of the RUC forecast errors at both waypoints. In the East direction, more than 50% of the forecast errors tend to map to positive error states; this behavior is not evident in the North direction, where the median error falls in the state centered on 0 m/s , s_6 .

2.3.3.6 Transition Probabilities

A further classification of forecast errors is required, however, to construct the inhomogeneous Markov model defined in section 2.3.1. Recall that the probability transition matrix $(P^{n,m})_{m \geq n}$ was previously defined:

Definition $(P^{n,m})_{i,j} = \sum_{\ell \in S} p_{i,\ell}^{n,(m-1)} p_{\ell,j}^m$




Figure 9: Normalized histogram of RUC forecast errors at MWH for the base case model.

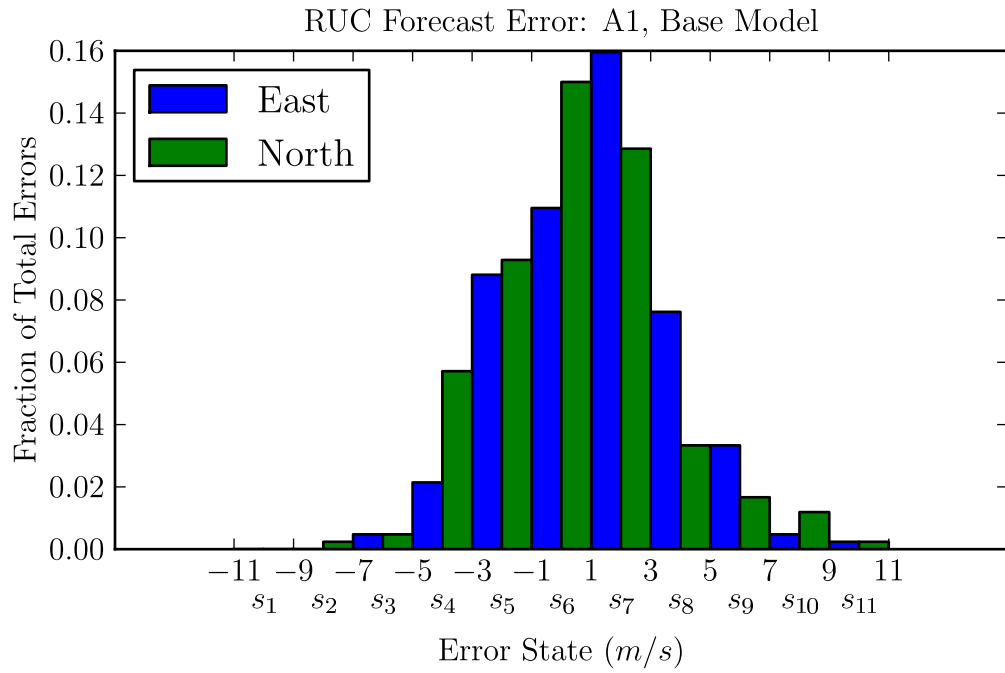


Figure 10: Normalized histogram of RUC forecast errors at A1 for the base case model.




Figure 11: Cumulative histogram of RUC forecast errors at MWH for the base case model.

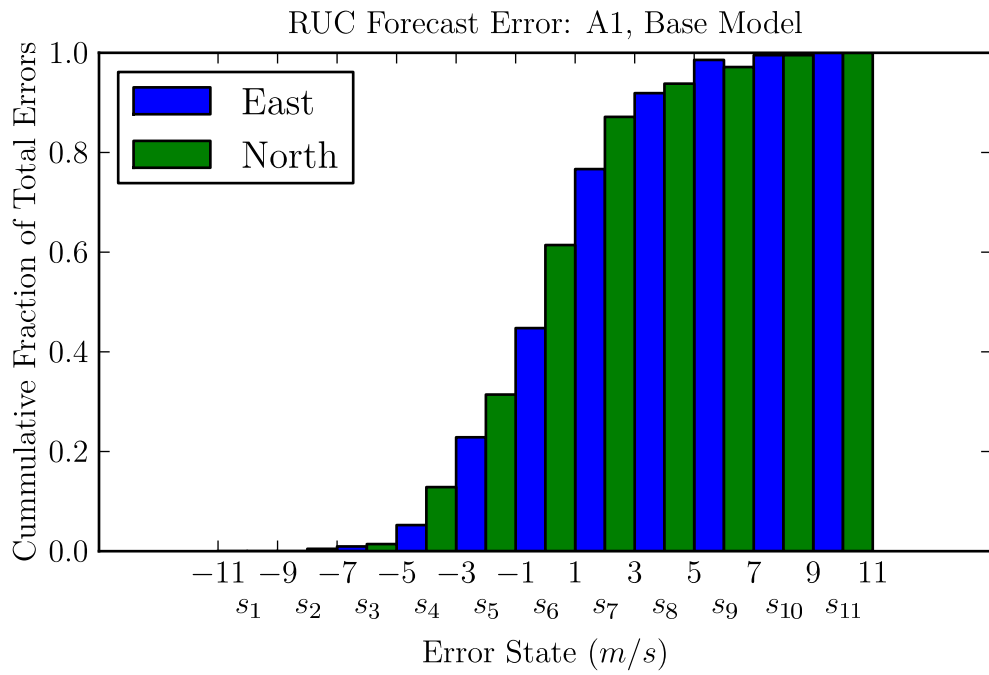


Figure 12: Cumulative histogram of RUC forecast errors at A1 for the base case model.

Definition 2.3.1 provides a means of calculating state transition probabilities between a location n and m given information is known about the state transition probabilities between points n and a point $m - 1$, which is only one step from m . For example, consider the sample route between KSEA and KATL. In order to calculate the state transition probabilities between waypoints MWH and A03, the route between the waypoints must be considered: MWH, A1, A2, A3. Given Definition 2.3.1, one must at the very least know the state transition probabilities between waypoints MWH and A2, as A2 is the only waypoint within one prior step of A3:

$$(P_{i,j}^{MWH,A3}) = \sum_{\ell \in S} P_{i,\ell}^{MWH,A2} p_{\ell,j}^{A3}$$

However, $(P_{i,j}^{MWH,A2})$ can similarly be constructed given information regarding the state transition probabilities between waypoints MWH and A1, as A1 is the only waypoint within one prior step of A2:

$$(P_{i,j}^{MWH,A2}) = \sum_{\ell \in S} P_{i,\ell}^{MWH,A1} p_{\ell,j}^{A2} \quad (11)$$

Equation 11 leaves the simplest case, a step between locations MWH and A1 requiring the state transition probability matrix $(P_{i,j}^{MWH,A1})$. Recall that a state probability transition matrix is simply a matrix indexed by indices on $S \times S$, where states $s \in S$. In this scenario, S is finite and s_i is indexed $1, 2, \dots, 11$, meaning $(P_{i,j}^{MWH,A1})$ is simply an 11×11 matrix tracking the state transition probabilities between locations MWH and A1. Furthermore, each row $(P_{i,*}^{MWH,A1})$ is a distribution λ_i describing the probability of the forecast error at A01 given the error state s_i at MWH. As the sample route explicitly requires the aircraft to proceed from MWH to A1 directly, calculating $(P_{i,j}^{MWH,A1})$ is an exercise in accounting of the error state transitions for individual flight segments between MWH and A1. Consider the illustration provided in Figure 13. The state transition probability $p_{1,1}^{1,2}$ describes the probability of a forecast error state transitioning from state s_1 to state s_1 as the aircraft moves from location (1) to location (2). Similarly, $p_{5,4}^{1,2}$ describes the probability of a forecast




Figure 13: Example of state transition accounting.

error state transitioning from state s_5 to state s_4 as the aircraft moves from location (1) to location (2). Calculating any $p_{i,j}^{1,2}$ for a one-step transition from location (1) to location (2) can be done directly, by finding the fraction of errors moving from a state s_i at location (1) to a state s_j at location (2) out of the total number of errors in state s_i at location (1). The same methodology applies to calculating the rows $(P^{MWH,A1})_i$. As the forecast errors are recorded as a function of individual ACARS flights, finding the fraction of forecast errors transitioning to a state s_i at MWH to s_j at A1 is again a matter of simple accounting of error transitions between the locations across the entire set of ACARS flights. As a reminder, each row $(P_{i,*}^{MWH,A1})$ is a distribution λ_i describing the probability of the forecast error taking a state s_j at A1 given the state s_i at MWH.

$(P^{MWH,A1})$ for the East component of the base case model is illustrated in Figure 14. Referring to the notation discussed above, each row of the transition probability matrix is indexed by a forecast error state at waypoint MWH; these error states are noted along the vertical axis of Figure 14. The corresponding forecast error states at waypoint A1 index the columns of $(P^{MWH,A1})$ and are listed along the horizontal axis

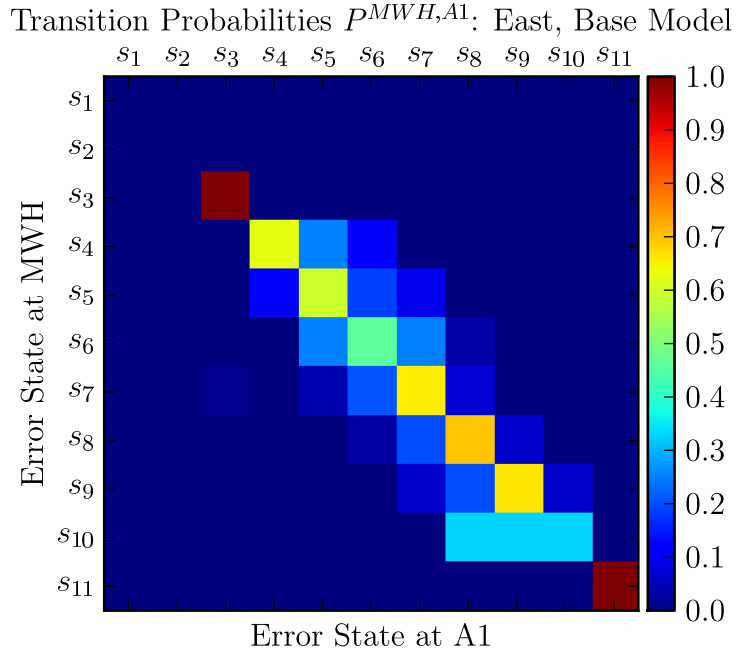


Figure 14: Illustration of transition probability matrix ($P^{MWH,A1}$) for the East component of the base case model.

of Figure 14. Moving from s_1 to s_{11} down the vertical axis, the rows of ($P^{MWH,A1}$) are distributions describing the probability of the forecast error taking a state s_j at A1. For example, if the forecast error is found to be in state s_7 at waypoint MWH, it is most likely to remain in state s_7 at waypoint A1 with a probability of just over 60%. The next most likely transition for an error in state s_7 at MWH is to transition to state s_6 at A1; this transition occurs with a probability near 25%. As a further note to the reader, areas of high probability near the edge states such as the $p_{3,3}^{MWH,A1}$ and $p_{11,11}^{MWH,A1}$ transitions are usually the result of fewer error records in those initial states rather than high probability of transition, specifically.

2.3.3.7 The Inhomogeneous Markov Chain

Now that the simplest case of a one step transition has been constructed, the entire inhomogeneous Markov chain can be constructed as a series of one step transitions along

each flight segment $\{\text{MWH-A1}, \text{A1-A2}, \dots, \text{A15-BNA}\}$ by calculating the appropriate one step transition probability matrices $\{(P^{MWH,A1}), (P^{A1,A1}), \dots, (P^{A15,BNA})\}$ and using Definition 2.3.1 as demonstrated in Equation 11. Referring to Section 2.3.1, however, one will find that a final piece of the puzzle is required in order to generate wind forecast scenarios using the model: an initial distribution. Recall that the definition of the model stated that a Markov chain $(X_n)_{n \geq 0}$ is defined as a stochastic process holding the following properties in terms of a stochastic transition probability matrix P and initial distribution λ :

1. $\mathbb{P}(X_0 = s_0) = \lambda_{s,0}$
2. $\mathbb{P}(X_{n+1} = s_{n+1} | X_0 = s_0, \dots, X_n = s_n) = p_{s_n, s_{n+1}}$

for $n \geq 0$ and every $s_0, \dots, s_n \in S$. $(X_n)_{0 \leq n \leq N}$ is *Markov* (λ, P) if $(X_n)_{0 \leq n \leq N}$ is a sequence of random variables satisfying (1) and (2) above for $n = 0, 1, \dots, N - 1$. In order to generate the model and subsequently sample realizations in order to form forecast error scenarios, some information regarding the initial distribution of the error states is required.

In a real-world scenario, actual wind speed error information may be used to specify the initial distribution λ . For example, consider an aircraft following the sample route from KSEA to KATL. The aircraft will have both a forecast of the wind speeds at its initial waypoint MWH as well as a measure of the actual winds at MWH, meaning an initial forecast error can be calculated. If the forecast error at MWH is measured to be 3.5 m/s , the initial error is known to exist in state s_8 . The initial distribution λ_{MWH} can therefore be defined to reflect this known information $\lambda_{MWH} = \{0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0\}$ (i.e.: the forecast error exists in state s_8 with 100% probability). The notional scenario posed to this point does not provide for an injection of known information into the model. Accordingly, the initial distribution at MWH is calculated directly; this distribution has already been presented in Figure

9.

Given an initial distribution, the inhomogeneous Markov chain is now simply calculated using the constructed one step transition probability matrices:

$$(X_n) = \lambda_0 \cdot \prod_{n=0}^N (P^{n,n+1}) \quad (12)$$

Product notation is used in Equation 12 for simplicity. However, the summation notation used in Definition 2.3.1 also applies. Again considering the scenario of an aircraft following the sample route from KSEA to KATL, Equation 12 is used to construct the Markov model sequentially as follows:

$$(X_{A1}) = \lambda_{MWH} \cdot (P^{MWH,A1}) \quad (13)$$

$$(X_{A2}) = \lambda_{MWH} \cdot (P^{MWH,A1})(P^{A1,A2}) = (X_{A1})(P^{A1,A2}) \quad (14)$$

$$(X_{BNA}) = \lambda_{MWH} \cdot (P^{MWH,A1})(P^{A1,A2}) \dots (P^{A15,BNA}) = (X_{A15})(P^{A15,BNA}) \quad (15)$$

The process described in Equations 13-15 is further illustrated in Figures 15 and 16 using the results of the forecast error analysis along the sample route. The evolution of the inhomogeneous Markov chain for the East component of the base case model is depicted in Figure 15, whereas the evolution of the inhomogeneous Markov chain for the North component of the model is illustrated in Figure 16. The vertical axes list the probable error states in ascending order from top to bottom, the horizontal axis tracks the age of the chain in terms of the route waypoint, and the intensity of the individual cell indexed (s_i, j) is the probability of the forecast error existing in state s_i at waypoint j . In practice, the model is better understood as a function of columns. Each column of cells in Figures 15 and 16 represents the distribution of the forecast errors λ_j at a particular waypoint j . For example, the columns index A4 in Figures 15 and 16 provide the distribution of forecast uncertainty λ_{A4} at waypoint A4 in the for the East and North components of the base case model respectively. Moving from left to right, one tracks the probability distributions of forecast uncertainty as a function

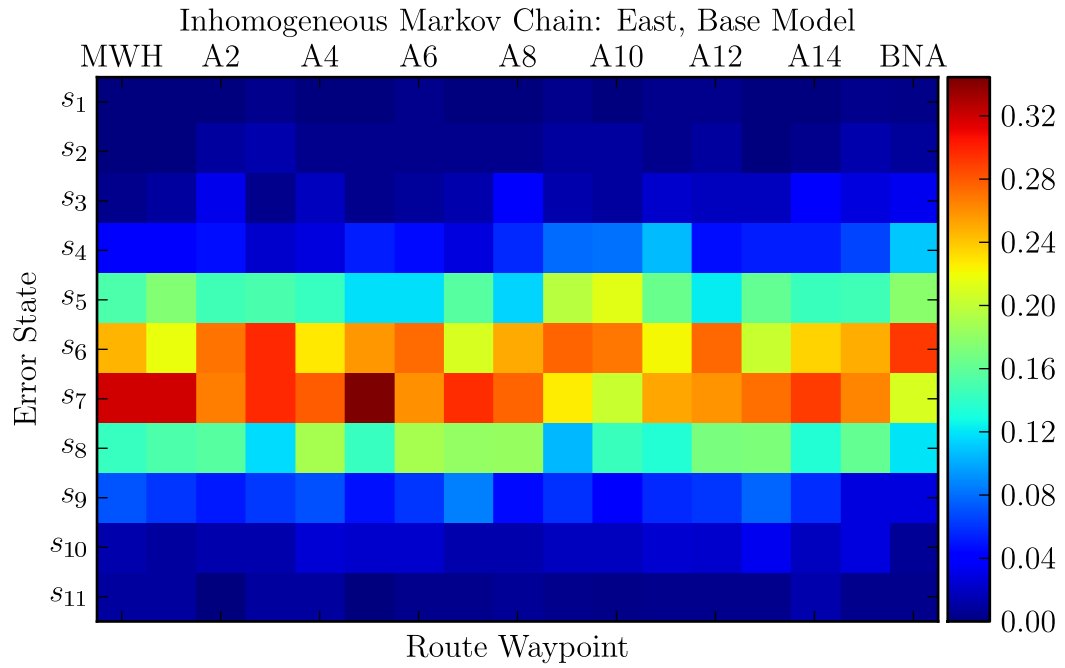


Figure 15: Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model.




Figure 16: Illustration of the calculated inhomogeneous Markov chain for the North component of the base case model.

of route progress. As an important reminder, based on the definition of the model derived in Section 2.3.1, the results presented in Figures 15 and 16 are functions of the specific series of waypoints visited in order along the sample route from KSEA to KATL.

Table 5: Wind forecast uncertainty model analysis routes.

Route Number	Origin Airport	Destination Airport
1	KSEA	KATL
2	KATL	KSEA
3	KLAX	KATL
4	KATL	KLAX
5	KEWR	KATL
6	KATL	KEWR

2.4 Selected Results

An inhomogeneous Markov approach to wind forecast uncertainty modeling was detailed in Section 2.3. A sample model was constructed using ACARS flights tracks along a sample route between Seattle’s Tacoma International Airport (KSEA) and Hartsfield-Jackson International Airport in Atlanta (KATL). In this section, the modeling methodology is applied to five more routes between Los Angeles International Airport (KLAX) and KATL, and Newark Liberty International airport (KEWR) and KATL. In total, six complete wind forecast uncertainty models between each of the three partner airports and KATL are considered; these models are defined in terms of origin and destination airport pair in Table 5. Detailed descriptions of these routes including waypoint lists and corresponding plots consistent with Table 2 and Figure 6 respectively are provided in Appendix A.

2.4.1 Jet Stream Effects

Jet streams are narrow bands of high velocity winds in the upper atmosphere which flow along the boundaries between hot and cold air masses. In the continental United States, jet streams flow from West to East but vary in daily position along a North-South axis [65]. The experiences of World War II fighter pilots who encountered strong westerly winds in both the European and Asian theaters of war (a summary of these documented encounters is available in [53]) led to seminal works by “Staff Members” of the University of Chicago’s Department of Meteorology [83], Rossby

[73], and Riehl [71] characterizing the high velocity atmospheric winds now referred to as jet streams. Though specifics regarding the meteorology behind the formations of jet streams are not important to this research effort, the usefulness of jet streams to aviation is of key interest.

Recall from Chapter 1 that an aircraft's ground speed is a function of both its planned airspeed and the wind speed it encounters along the route. An aircraft encountering a strong tailwind assisting its motion will travel at a considerably higher groundspeed than an aircraft flying into a headwind directly impeding its motion, resulting in a shorter flight time and lower fuel expenditure for the same planned airspeed. Clearly it is advantageous for commercial carriers to plan their routes in such a manner as to leverage the tail wind effects of the jet stream for flights from West to East, while avoiding the headwind effects of the jet stream for routes from East to West. Now consider that aviation wind speed reports constitute a large portion of the baseline data used to generate the RUC forecast (discussed in Section 2.1.1.1). The result of wind-favorable route planning by commercial air carriers is that the RUC forecast is initialized using a high percentage of wind speed measurements from within the jet stream.

Figures 17 and 18 depict the inhomogeneous Markov chains in the East direction for the base model for routes 1 (KSEA to KATL) and 3 (KLAX to KATL) respectively.

Routes 1 and 3 are eastward routes, along which aircraft are likely to leverage tailwinds provided by the jet stream. In both Markov chains, there is a tendency of the forecast error to skew towards error state s_7 along the entire route, where state s_7 is on the range between 1 $m/$ and 3 m/s . In other words, the forecast error is more likely to be slightly positive biased than zero-centered or negative, meaning the RUC forecast slightly over predicts the wind speeds along these routes. Similar behavior is seen along route 6 from KATL to KEWR in Figures 19 and 20. In this case, it is common for the jet stream to turn from East to North and follow the eastern sea




Figure 17: Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model for the KSEA to KATL route.

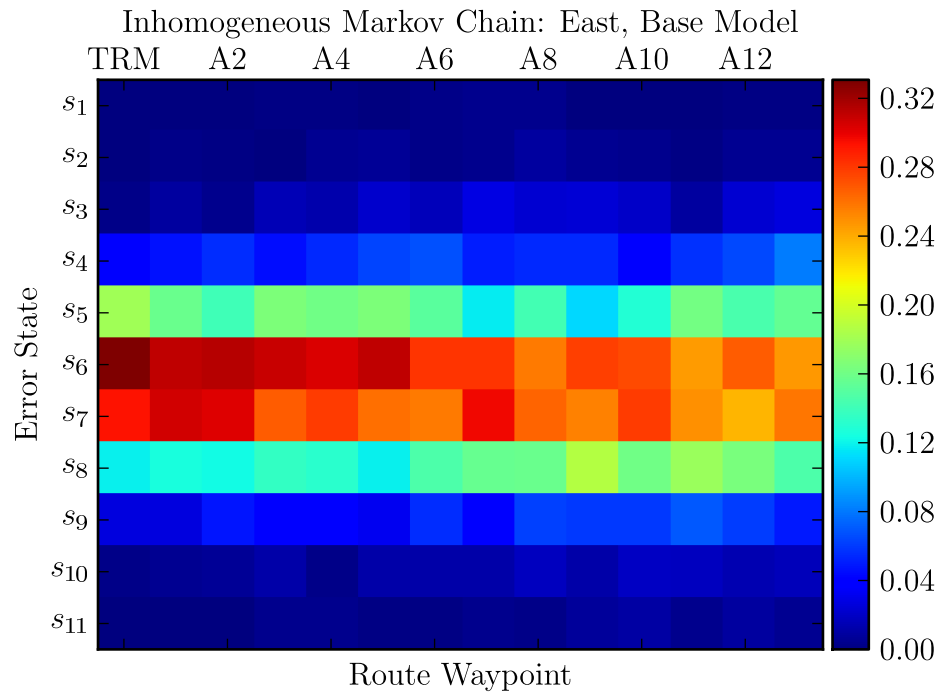


Figure 18: Illustration of the calculated inhomogeneous Markov chain for the North component of the base case model for the KLAX to KATL route.

Inhomogeneous Markov Chain: East, Base Model




Figure 19: Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model for the KATL to KEWR route.

Inhomogeneous Markov Chain: North, Base Model




Figure 20: Illustration of the calculated inhomogeneous Markov chain for the North component of the base case model for the KATL to KEWR route.

board of the United States, essentially flowing along route 6. A heavy Eastern bias is seen in Figure 19 at the beginning of the route near the jet stream turning point, the effects of which linger along the rest of the route. A consistent northern bias is seen in Figure 20, which mimics the results seen along routes 1 and 3 depicting a consistent over prediction by the RUC in the direction of the jet stream.

These results reinforce prior findings by the author [86] which found that the aviation reports tend to add a slight positive bias in favor of the jet stream due to wind-optimal route planning by air carriers. Conversely, the models for routes opposing the direction of the jet stream (routes 2, 4, and 5) exhibit a slight negative bias, meaning the RUC under predicts the wind speeds in regions not affected by a jet stream; these results are found in Figures 21 and 22. In both models, the error along the route has a tendency to skew towards error state s_5 defined on the range between $(-1) m/s$ and $(-3) m/s$. Though the positive jet stream bias of the RUC was previously confirmed by the author in [86], the tendency of the RUC to slightly under predict wind speeds along routes opposing the direction of the jet stream is a new finding.

2.4.2 Updated Forecasts

Section 2.3.2.3 discussed two possible paths through the RUC forecast time horizon corresponding to “base” and “updated” cases. The base case considered the forecast release to be fixed throughout the analysis period whereas the updated case only used the freshest 1-hour forecast release available. Figures 23 and 24 depict the Markov models for the East components of route 4 between KATL and KLAX for both the base and updated cases. Though slight, there is a difference in the base and updated model. The updated model depicted in Figure 24 depicts slightly narrower probability bounds in the later stages of the route. For example, consider the outer edges of the forecast uncertainty distributions in both models, specifically states s_3 and s_8 . The




Figure 21: Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model for the KATL to KSEA route.




Figure 22: Illustration of the calculated inhomogeneous Markov chain for the North component of the base case model for the KATL to KLAX route.

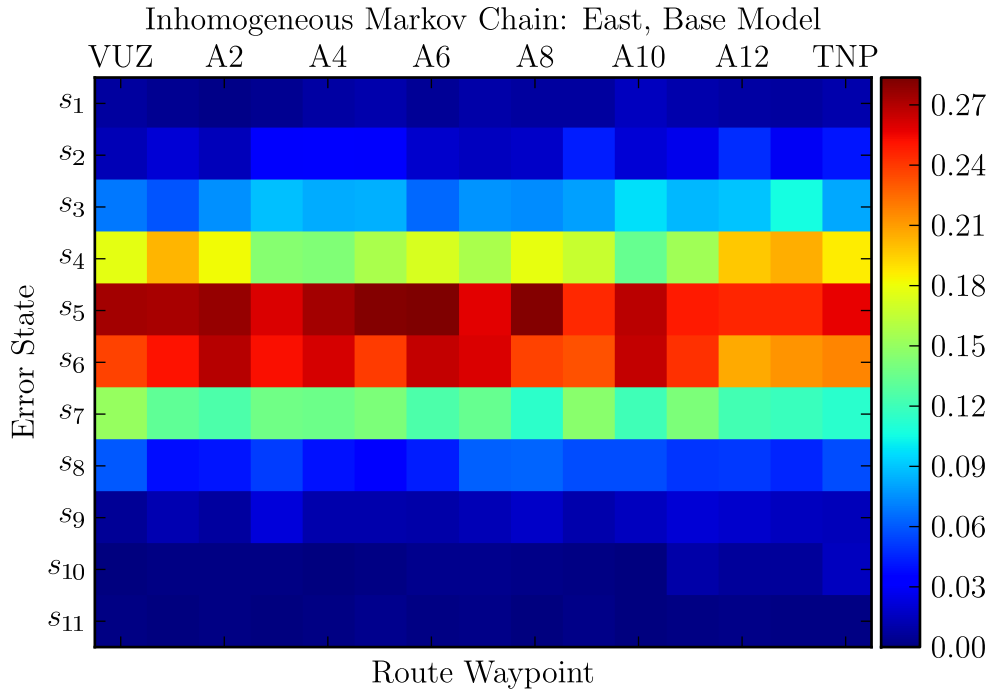


Figure 23: Illustration of the calculated inhomogeneous Markov chain for the East component of the base case model for the KATL to KLAX route.

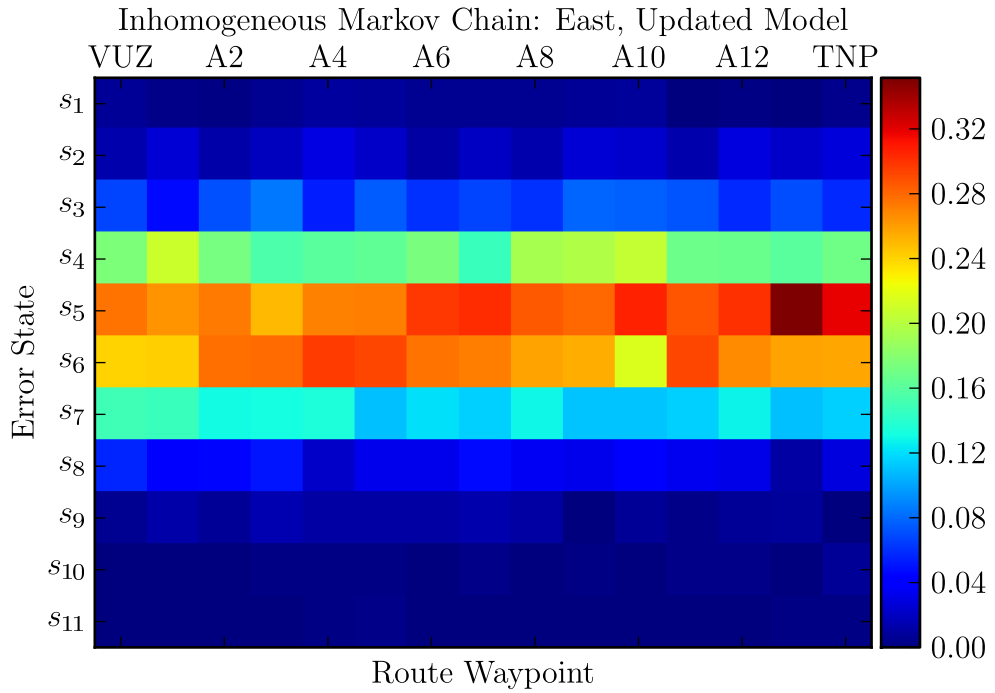


Figure 24: Illustration of the calculated inhomogeneous Markov chain for the East component of the updated case model for the KATL to KLAX route.

base model depicted in Figure 22 demonstrates a higher likelihood of the forecast error existing in these outer states, as one moves from VUZ to TNP. Though this result is again very slight, it does lend credence to the idea that updating wind information along the route does slightly improve forecast skill as demonstrated by a decreased forecast uncertainty. Though only the result for route 4 have been provided, this result is consistent across the route set.

CHAPTER III

THE RTA PROBLEM

A review of forecast uncertainty modeling literature, a detailed exploration of the forecast uncertainty modeling methodology utilized in this research effort, and selected results along sample routes between sets of origin and destination airports were presented in Chapter 2. This chapter switches focus from uncertainty modeling to stochastic optimization, specifically as it relates to the fuel optimal required time of arrival (RTA) problem introduced in Chapter 1. An introduction to stochastic programming methodology is first presented. The unique challenges of the fuel optimal RTA problem are then discussed before the research gaps in this area are identified. The solution algorithm is then developed before selected results for flight scenarios along the sample routes constructed in Chapter 2 are presented and discussed.

3.1 Stochastic Programming

3.1.1 A Simple Recourse Model

Stochastic models describe systems with uncertain inputs that can be characterized by known probability distributions. As opposed to deterministic models which assume the input data is known at time of solution, stochastic models are used to account for variability in possible scenarios given the realization of data uncertainty as time progresses. The most common stochastic model is the recourse model. In a general recourse model, a decision is made at the current point in time before any uncertainty has been realized. The goal of this decision is to minimize the net cost of the current decision and any future necessary corrective decisions as the uncertainty scenarios unfold. The simplest recourse model is the two-stage model which relies on a single recourse decision. An initial decision is made for the first stage, the uncertainty in

the data is realized, and a recourse decision is made for the second stage. In the context of the fuel optimal RTA problem, consider the forecasted wind speeds along the route to be the source of uncertainty. The task is to make a decision on the initial chosen cruise speed, given this cruise speed may need to be changed at some point in the future in order to satisfy a given RTA. The goal of this initial cruise speed is to minimize the total expected fuel burn given all possible correction scenarios required due to arrival time estimation error that results from wind speed forecast uncertainty.

A linear two-stage model as introduced by Beale [1] and Dantzig [21] is presented below:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & c^\top x + \mathbb{E}_\xi [Q(x, \xi)] \\ \text{s.t.} \quad & Ax = b \\ & x \geq 0 \end{aligned} \tag{16}$$

where x is the first stage decision vector, $c^\top x$ is the cost of the first stage decision, and $\xi(q, T, W, h)$ is the data associated with the second stage model. In the context of the *linearized* RTA problem, x contains choices for the initial cruise speed, $c^\top x$ contains the fuel costs associated with selecting the initial cruise speeds, and $\xi(q, T, W, h)$ contains data for the second stage model including wind speed scenarios generated using the forecast uncertainty model. The solution to the second stage model is defined:

$$\begin{aligned} Q(x, \xi) = \min_{y \in \mathbb{R}^m} \quad & q^\top y \\ \text{s.t.} \quad & Tx + Wy = h \\ & y \geq 0 \end{aligned} \tag{17}$$

In order to solve the two-stage model numerically, it can be assumed that the random vector ξ contains a finite number of possible realizations, referred to as scenarios $(\xi_1, \xi_2, \dots, \xi_K)$ with probabilities of occurrence (p_1, p_2, \dots, p_K) . As the expectation of the second stage solution can then be expressed as a function of its scenarios:

$$\mathbb{E}_\xi [Q(x, \xi)] = \sum_{k=1}^K p_k Q(x, \xi_k) \tag{18}$$

the two stage model presented in Equations 16-17 can be formulated as one large-scale linear programming model:

$$\begin{aligned}
& \min_{x, y_1, y_2, \dots, y_k} && c^\top x + \sum_{k=1}^K p_k q_k^\top y_k \\
& s.t. && Ax = b \\
& && Tx_k + Wy^k = h_k \quad \forall k \in (1, 2, \dots, K) \\
& && x \geq 0 \\
& && y_k \geq 0
\end{aligned} \tag{19}$$

Cutting plane methods are a common solution strategy to solving large-scale linear programs. Cutting plane methods redefine the model’s feasible space by iteratively adding linear inequality constraints to the model. The first cutting plane method was proposed by Kelley in 1960 [44], and formed the basis of later methods including both Dantzig-Wolfe [20, 19] and Bender’s [2] decompositions methods. Decomposition methods segregate the problem into a first stage master problem based on the first stage decision and multiple second stage problems defined by the realizations of the random vector ξ described in Equations 16 - 19. Several more cutting plane and decomposition methodologies have been developed to improve computational efficiency including the multi-cut method proposed by Birge and Louveaux [10] and the regularized decomposition method proposed by Ruszczyński [75]. These works will not be discussed in detail, but are provided as reference for the reader.

3.1.2 Relationship to Dynamic Programming

In developing an algorithm to explore the solution space posed by a stochastic programming problem, it is helpful to draw parallels between the classic two-stage linear stochastic programming model introduced in Equations 16 - 19 and dynamic programming approaches. Recall again that in a stochastic programming framework, the uncertainty data $(\xi_1, \xi_2, \dots, \xi_K)$ is revealed over a time horizon, and that in the setting of a stochastic decision process, decisions must be made before realizations of

the model uncertainty come to pass. The decision process therefore has the form:

$$\begin{aligned} & \text{decision}(x_1) \rightarrow \text{observation}(\xi_1) \rightarrow \text{decision}(x_2) \rightarrow \\ & \dots \rightarrow \text{decision}(x_T) \rightarrow \text{observation}(\xi_T) \end{aligned} \quad (20)$$

where the sequence $\xi_t \in \mathbb{R}^{d_t}$ for $t = 1, 2, \dots, T$ of data vectors is a stochastic process with known probability distributions (note that $\xi_{[t]} = (\xi_1, \xi_2, \dots, \xi_t)$ refers to the history of the process up to time t). The form of the decision process described in Equation 20 introduces the concept of nonanticipativity. Again, assume that the random vector ξ contains a finite number of possible realizations or scenarios $(\xi_1, \xi_2, \dots, \xi_K)$ with probabilities of occurrence (p_1, p_2, \dots, p_K) . The two-stage formulation presented in Equation 16 can be relaxed by replacing the decision vector x with K vectors (x_1, x_2, \dots, x_K) , one for each scenario. This relaxation results in K smaller problems of the form:

$$\begin{aligned} & \min_{x_k \geq 0, y_k \geq 0} \quad c^\top x_k + q_k^\top y_k \\ & \text{s.t.} \quad Ax_k = b \\ & \quad \quad T_k x_k + W_k y_k = h_k \end{aligned} \quad (21)$$

Based on the decision process presented in Equation 20, it is clear that the model posed in Equation 21 is unsuitable for finding a two-stage decision process as the first stage decision variables x_k depend on realizations of the random data from the second stage. An additional constraint, referred to as a nonanticipativity constraint, is required to ensure that the first stage decision variables are independent of the realizations of the random data. One of the simplest ways to enforce nonanticipativity is to require that all realizations of the first stage decision vector are equal to one another. Another is to set all realizations of the first stage decision vector equal to the weighted average of the entire set of first stage scenarios weighted by probability of occurrence. Both implementations ensure the requirement of nonanticipativity, enforcing that x_t depend on ξ_t alone, and not future observations.

Given the problem must be formulated to reflect nonanticipativity, consider a mutli-stage problem with T stages, specifically the final-stage problem at period T :

$$\min_{x_T \in \mathcal{X}_T(x_{T-1}, \xi_T)} f_T(x_T, \xi_T)$$

where $x_t \in \mathbb{R}^{n_t}$ are decision variables, f_t are defined on $\mathbb{R}^{n_t} \times \mathbb{R}^{d_t}$, and \mathcal{X}_t are defined on $\mathbb{R}^{n_{t-1}} \times \mathbb{R}^{d_t}$ for $t = 1, 2, \dots, T$. Recall based on the structure of the decision process that the optimal value of this problem, denoted $Q_T(x_{T-1}, \xi_T)$, is a function of the prior stage decision vector x_{T-1} and current stage data ξ_T . At stage $T - 1$, the problem becomes:

$$\min_{x_{T-1} \in \mathcal{X}_{T-1}(x_{T-2}, \xi_{T-1})} f_{T-1}(x_{T-1}, \xi_{T-1}) + \mathbb{E} [Q_T(x_{T-1}, \xi_{[T]}) \mid \xi_{[T-1]}]$$

where $\mathbb{E} [\cdot \mid \xi_{[t]}]$ is a conditional expectation, meaning the $T - 1$ stage's optimal value depends on the optimal value of stage T , $Q_T(x_{T-1}, \xi_T)$, given the realization of the uncertainty process ξ at stage $T - 1$. This trend continues as one steps through prior states, allowing prior stage formulations to be written more generally for stages $t \in 2, 3, \dots, T$:

$$\min_{x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)} f_t(x_t, \xi_t) + \mathbb{E} [Q_{t+1}(x_t, \xi_{[t+1]}) \mid \xi_{[t]}]$$

As a result, solving the any single stage requires recursively calculating prior stage “cost-to-go” functions $Q_t(x_{t-1}, \xi_{[t]})$, as they are functions of the prior stage decisions x_{t-1} and the realizations of the uncertainty data $\xi_{[t]}$. Eventually, the entire time history must be explored, leading to the first stage problem:

$$\min_{x_1 \in \mathcal{X}_1} f_1(x_1) + \mathbb{E} [Q_2(x_1, \xi_{[2]})]$$

The equivalent dynamic programming equations are written [80]:

$$Q_t(x_{t-1}, \xi_{[t]}) = \inf_{x_t \in \mathcal{X}_t(x_{t-1}, \xi_t)} [f_t(x_t, \xi_t) + Q_{t+1}(x_t, \xi_{[t]})] \quad (22)$$

where

$$Q_{t+1}(x_t, \xi_{[t]}) := \mathbb{E} [Q_{t+1}(x_t, \xi_{[t+1]}) \mid \xi_{[t]}] \quad (23)$$

Given this construct, the decision variables $x_t = \mathbf{x}_t(\xi_{[t]})$ can be considered to be functions of the data process $\xi_{[t]}$ up to time t . A sequence of feasible decision variables representing a series of feasible decisions to be made at each stage is referred to as an “implementable policy.” An implementable policy is considered optimal if and only if for every $t \in 1, 2, \dots, T$,

$$x_t^*(\xi_{[t]}) \in \underset{x_t \in \mathcal{X}_t(\bar{\mathbf{x}}_{t-1}(\xi_{[t-1]}), \xi_t)}{\operatorname{argmin}} [f_t(x_t, \xi_t) + \mathcal{Q}_{t+1}(x_t, \xi_{[t]})]$$

with probability equal to 1 (note that for $t = T$, \mathcal{Q}_{t+1} is omitted; the set \mathcal{X}_t only depends on ξ_1 for $t = 1$). In other words, an optimal implementable policy provides a sequence of optimal stage-wise decisions over the evolution of the problem uncertainty $\xi_{[t]}$ (the reader should be aware that an optimal implementable policy considers future optimal recourse in prior stage decisions based on the nested nature of the prior stage formulation). Essentially, a dynamic programming approach reduces the problem to a series of finite dimensional problems indexed by t and $\xi_{[t]}$.

3.1.3 Scenario Trees

Recall that the random process ξ contains a finite number of possible realizations or scenarios $(\xi_1, \xi_2, \dots, \xi_K)$ with probabilities of occurrence (p_1, p_2, \dots, p_K) . The process of constructing these scenarios in such a manner as to both capture the structure of the proposed decision process as well as accurately discretizing the forecast uncertainty process is the next area of focus.

Scenario trees are useful tools for visualizing the evolution of uncertainty processes over multiple problems stages. In general, a scenario tree is constructed of several levels, organized by problem stages $1, 2, \dots, T$, each layer containing multiple nodes. The first level, denoted by $t = 1$, contains only a “root” node, which has a known associated value ξ_1 . Subsequent levels $t + 1$ contain as many nodes ι as there are possible realizations of ξ_{t+1} , referred to as “child” nodes, where each child node is




Figure 25: Sample scenario tree.

connected to an “ancestor node” in level $t - 1$ via an arc. Based on the parent-child relationship of the nodes between each level, all of the nodes in a given level t represent possible realizations of the process ξ_t , and share common ancestry of the process history $\xi_{[t]}$. Thus, a “scenario” as described to this point, is a path from the root node through the levels of a scenario tree, constituted by a series of individual realizations ξ_t^i , where t is again the problem stage, and i a specific node within level t . Ω_t denotes the set of all nodes in a given stage $t \in 1, 2, \dots, T$. A scenario tree with $t = 3$ levels and varying number of nodes i in each level is depicted in Figure 25.

Given the nested nature of the decision process defined in Equation 20, in order to define probability distributions describing the uncertainty process, the process must be modeled using conditional distributions. Specifically, the process ξ_{t+1} is a function of the specific time history of the process $\xi_{[t]}$, $t = 1, 2, \dots, T - 1$. If the uncertainty process holds a value b_i at a node $i \in \Omega_t$, one can specify the probability of the process holding a value b_j at a node $j \in \Omega_{t+1}$. Let this probability be defined as $p_{i,j}$, and lie along the arc connecting nodes i and j . Furthermore, require that $p_{i,j} \geq 0 \forall i, j \in \Omega_t$, and $\sum_{j \in \Omega_{t+1}} p_{i,j} = 1$. Referring again to Figure 25, one can see that the uncertainty process holds value b_{l_1} at stage $t = 1$ with probability equal




Figure 26: Sample decision scenario for route 1 between KSEA and KATL.

to 1 (as it is the root node); the process will hold value b_{l_1} at stage $t = 2$ (distinct from b_{l_1} with $t = 1$) with probability p_{l_1, l_1} , and value b_{l_2} with probability p_{l_1, l_2} . It should be apparent to the reader that the probabilities along any given stage t form a conditional distribution describing the state of the process ξ at stage t , given the process $\xi_{[t-1]}$.

The inhomogenous Markov model developed in Chapter 2 is well-conditioned to generating wind forecast uncertainty scenarios using a node-based method. Again consider the decision process defined in Equation 12; however, this time, consider it in the context of an aircraft cruising along one of the sample routes previously discussed, specifically route 3 between KLAX and KATL (described in Appendix A.2). Referring to Figure 26 and Equation 20, the decision process dictates that the aircraft make an initial choice in Mach number M_1 at waypoint MWH in order to minimize the total fuel burn along the route given this choice will later be revised to M_2 at artificial waypoint A10 to correct for flight time estimate errors introduced by wind forecast uncertainty. Assume the wind forecast uncertainty at MWH, ξ_{MWH} is known; ξ_{MWH} holds one value which maps to a known error state $s_{MWH} \in S$, and serves as the root node of the scenario tree (a singleton distribution can be used to describe the probability distribution of a known value in the Markov model). The possible error states at A10 s_j with $j = 1, 2, \dots, 11$ each map to a child node ι_{s_j} , which connect to the root node via an arc with probability $p_{s_{MWH}, \iota_{s_j}}^{MWH, A10}$. Recall from Section 2.3.3.5 that the set of probabilities $p_{s_{MWH}, \iota_{s_j}}^{MWH, A10}$ constitute the conditional distribution λ_{A10} ; in the context of the scenario tree description, λ_{A10} describes the uncertainty process $\xi_{[A10]}$ given the process's known state at MWH, s_{MWH} . Given $\xi_{[A10]}$, subsequent tree

levels are similarly constructed. The possible error states at BNA map to child nodes η_{s_j} , each of which connects to a corresponding parent node in the MWH level via an arc with probability $p_{\iota_{s_j}, \eta_{s_j}}^{A10, BNA}$. Similarly, λ_{BNA} can be constructed to describe $\xi_{[BNA]}$, which is a function of both $\xi_{[A10]}$ and ξ_{MWH}

The Markov property was introduced in Section 2.3.1. In the simplest terms, a stochastic process was said to be Markovian if it held the Markov property, meaning that its future state was a function of its current state alone and not the time history of prior states. The scenario tree construction presented above constitutes the entire time history of the process whether or not the process hold the Markov property, as every possible state combination is enumerated as one steps through the scenario tree levels. Formulating an optimization problem from a dynamic programming standpoint as posed in Equations 22 - 23 eliminates the requirements to track the entire time history of the uncertainty process, as the stage cost-to-go functions $Q_t(x_{t-1}, \xi_{[t]})$ are functions of the current state of the uncertainty process $\xi_{[t]}$ alone. In essence, a dynamic programming approach leverages the Markovian structure of the process to reduce the information required at any one time. A stochastic programming approach conversely considers the entire time history of the process, thus making it more general at the expense of a requiring a significantly larger scenario set. This difference in scenario structure is the basic distinction between stochastic programming and dynamic programming approaches [80]. The challenge of a stochastic programming approach as it relates to this effort, is formulating an algorithm capable of exploring the entire scenario tree while maintaining computational tractability for onboard systems. Though a purely dynamic programming approach solves this particular problem, it does so at the expense of generality.

3.1.4 Nonlinearities

The model presented in Equations 16 - 19 assumes a linear relationship between the decision variables in x for both the objective function and the associated constraints. However, the fuel optimal RTA problem provides non-linearity in two areas: the relationship between travel distance, velocity, and time; and the relationship between an aircraft's fuel flow rate and velocity.

Referring to Equation 3, the distance travelled vs. flight time relationship is non-linear. Lowther provides a framework for linearizing the time-distance relationship by non-dimensionalizing velocity terms in terms of Mach number [56]. Recall that the distance flown by an aircraft in any given time period $t_i - t_0$ is represented by the integral equation:

$$d = \int_{t_0}^{t_1} \mathbf{v}_g \cdot \Delta t$$

Consider that the aircraft's ETA to the RTA waypoint differs by some value $\Delta t = t_{RTA} - t_{ETA}$. A constant speed change $\Delta \mathbf{v}_g$ can be introduced to achieve the change in arrival time Δt as the distance traversed by the aircraft remains the same (fixed path constraint) regardless of the flight speed:

$$\begin{aligned} d_{RTA} &= \int_{t_0}^{t_{ETA}-\Delta t} (\mathbf{v}_g + \Delta \mathbf{v}_g) dt \\ &= \int_{t_0}^{t_{ETA}} (\mathbf{v}_g + \Delta \mathbf{v}_g) - \int_{t_{ETA}-\Delta t}^{t_{ETA}} (\mathbf{v}_g + \Delta \mathbf{v}_g) \\ &= \int_{t_0}^{t_{ETA}} \mathbf{v}_g \cdot dt + \int_{t_0}^{t_{ETA}} \Delta \mathbf{v}_g \cdot dt - \int_{t_{ETA}-\Delta t}^{t_{ETA}} \mathbf{v}_g \cdot dt - \int_{t_{ETA}-\Delta t}^{t_{ETA}} \Delta \mathbf{v}_g \cdot dt \\ &= d_{RTA} + \Delta \mathbf{v}_g \cdot t_{ETA} - \int_{t_{ETA}-\Delta t}^{t_{ETA}} \mathbf{v}_g - \Delta \mathbf{v}_g \cdot \Delta t \end{aligned}$$

If it is further assumed that the groundspeed \mathbf{v}_g during time period $[t_{ETA} - \Delta t, t_{ETA}]$ remains constant such that $\mathbf{v}_g = \mathbf{v}_{g,i}$ (the groundspeed for this segment equals the initial groundspeed):

$$\Delta \mathbf{v}_g \cdot t_{ETA} - \mathbf{v}_{g,i} \cdot \Delta t - \Delta \mathbf{v}_g \cdot \Delta t = 0$$

a relationship between the arrival time change required to achieve an RTA and the

change in the current groundspeed required to affect that change can be defined:

$$t_{ETA} = \frac{(\mathbf{v}_{g,i} + \Delta\mathbf{v}_g) \cdot \Delta t}{\Delta\mathbf{v}_g}$$

Again recall that the groundspeed of the aircraft is a function of both the true air speed and the wind speeds encountered by the aircraft as described in Equation 2. If it is assumed that the cross-wind component (the component of the wind vector laying perpendicular to the aircraft's ground track) is relatively small compared to the aircraft's true airspeed, it can be said that the change in groundspeed required to satisfy a change in an ETA is approximately equal to an equivalent change in the aircraft's true airspeed. More specifically, a change in the aircraft's true airspeed can be introduced to affect an equivalent change in the aircraft's ground speed subject to the prevailing winds. Accordingly, Equation 3.1.4 is reformulated in terms of Mach number:

$$t_{ETA} = \frac{(M_i + \Delta M) \cdot \Delta t}{\Delta M} \quad (24)$$

where M is the Mach coefficient, the ratio of the aircraft's true airspeed \mathbf{v}_{TAS} and the local speed of sound a :

$$M = \frac{|\mathbf{v}_{TAS}|}{a}$$

Though now closer to forming a linear relationship between travel distance, true airspeed, and time, Equation 24 still provides bilinear unknown $\Delta M \cdot \Delta t$ when expanded via Taylor series. The only means of mitigating the affect of these nonlinear terms is to make a further assumption that ΔM is comparatively small to M_i . Lowther found that assuming ΔM has a maximum value of 0.02 introduced as much as 2.5% error into the flight time calculation for an M_i value of 0.8 [56]. This assumption is undesirable for two reasons: a speed change limit on Mach number severely reduces control authority, and increasing the speed change limit to a more acceptable value introduces considerable flight time estimation error as demonstrated by Lowther. Diminished control authority could conceivably prevent an aircraft from

satisfying an RTA if the speed change required to do so is beyond the limits imposed by the linearized model. Conversely, relaxing the speed change limit and consequently increasing flight time estimation error beyond the 2.5% margin demonstrated by Lowther is not an acceptable alternative given the adherence capabilities of current RTA systems.

The aircraft’s fuel burn characteristics introduce a second source of nonlinearity. Fuel burn rate is approximately quadratic as a function of Mach number [Boeing Performance Engineer’s Manual: 737-300/CFM56-3B-2]. Again, steps can be taken to linearize this relationship, this time by representing the quadratic relationship as a piecewise linear function. However, this methodology introduces the same ΔM limitation as the previous derivation in order to mitigate the affects of linearization and thus limits the available control authority as well.

A third, yet slightly more significant source of non-linearity is the forecast information itself. Lowther’s problem assumed a static wind forecast along the entire route for the duration of the flight. Regardless of the flight duration, the wind forecast remained fixed. The model presented in Chapter 2 posed two distinct yet still dynamic scenarios (referred to as the “base” and “updated” cases). In the base case, only wind information from a single release is used, whereas the updated model case considers only the first hour forecast of the newest freshest forecast release. In both cases, the wind forecast along the route updates on an hourly cycle (the distinction between the two cases is the source of the updated information). As the time versus distance travelled relationship is a function of the aircraft’s groundspeed, and the aircraft’s groundspeed is directly a function of the encountered winds along the route, the problem data itself changes as a function of the decision variables. In essence, the problem itself is purely dynamic. As a result, a linearized formulation as presented in Equation 19 is infeasible without completely neglecting the time component.

As one can see, formulating the fuel optimal RTA problem in terms of large-scale

linear program as posed in Equation 19 is undesirable due to linearization effects and the problem structure. Linearizing the distance travel versus flight time relationship introduces enough error to ensure that a linearized formulation will not perform to the same tolerances as currently available technologies while significantly limiting the control authority of the algorithm. More importantly though, a linearized model is incapable of capturing the time dynamics of the problem. The task is therefore to develop an algorithm that can address both the non-linearities inherent to the problem while maintaining computational tractability on an airborne system given the large solution space posed by a stochastic programming formulation.

3.2 Research Gaps

Current on-board systems fail to proactively address flight time estimation error introduced by wind forecast uncertainty as it relates to required time of arrival (RTA) capability for two primary reasons: the lack of a robust wind forecast uncertainty model, and the lack of a computationally tractable RTA algorithm capable of considering wind forecast uncertainty for speed planning purposes. A position-based wind forecast uncertainty model was developed in Chapter 2. To the second need, the shortcomings associated with linearizing the speed planning problem as a means of formulating the underlying stochastic program as a large-scale linear program were discussed in Section 3.1.4. There is a clear need for a computationally tractable RTA algorithm capable of being used by on-board systems to make speed planning decisions in the presence of wind forecast uncertainty that does not incur the penalties associated with a traditional approach utilizing linearization.

3.3 Algorithm Formulation

The simplest solution to the RTA problem as posed is a brute force enumeration of all possible scenarios across a set of possible cruise Mach numbers. However, as introduced in Section 3.1.3, this methodology is not computationally tractable in an on-board system given the limited computational capabilities of contemporary flight management systems. The development of a search heuristic capable of quickly exploring the entire set of wind speed forecast uncertainty scenarios is the next focus of this research effort, and the resulting algorithm the second contribution¹.

Similarly to the development of the forecast uncertainty model in Chapter 2, development of the RTA algorithm will rely on a sample scenario for illustrative purposes. Again, consider a flight between KSEA and KATL, this time, with some required time of arrival at the route’s final waypoint BNA. The route details are available in Section 2.3.3, and a description of this particular scenario was previously introduced in Figure 26. Assume that artificial waypoint A10 was chosen as the problem stage dividing point, meaning a recourse speed change will be made in order to account for wind forecast uncertainty.

3.3.1 Scenario Generation

Recall that the uncertainty model is initialized via an initial distribution. In the posed RTA scenario, this initial distribution reflects known forecast error information calculated by the aircraft as it reaches cruising altitude at MWH. As a reminder to the reader, the Markov model developed in Chapter 2 defined a series of $q = 11$ error states $s \in S$ on the range $R = [-11, 11]$, where S contains the error state intervals $\{(-11, -9], (-9, -7], \dots, (9, 11]\}$ indexed by $i = 1, 2, \dots, 11$. For example, if the forecast error at MWH is measured to be 3.5 m/s , the initial error is known to

¹The RTA algorithm is currently under review by the United States Patent and Trademark Office. Information regarding intellectual property claims can be found in [85].

exist in state s_8 . The initial distribution λ_{MWH} can therefore be defined to reflect this known information $\lambda_{MWH} = \{0, 0, 0, 0, 0, 0, 0, 1, 0, 0, 0\}$ (i.e.: the forecast error exists in state s_8 with 100% probability). In the context of the scenario tree as introduced in Section 3.1.3, λ_{MWH} provides a known value for the root node of the forecast uncertainty process ξ .

Given the known value of the uncertainty process at MWH, the conditional probability distribution of the uncertainty process ξ at the intermediate decision point A10 is constructed. This construction may be carried out in one of two ways: a transition probability matrix ($P^{MWH,A10}$) is calculated directly using error comparisons for data points neighboring MWH and A10, or ($P^{MWH,A10}$) is calculated indirectly using one-step transition probability matrices for intermediate waypoints between MWH and A10 (see Equation 12). Both methods have positive and negative consequences from a modeling standpoint. The first method is structurally sound from a scenario generation standpoint; a directly calculated conditional distribution reflects strictly the evolution of the process from MWH to A10, whereas the indirect method utilizing Equation 12 relies on intermediate history of the process to generate the conditional distribution at the next level. Utilizing the second methodology could require introducing a new scenario tree level as each intermediate conditional distribution is used to move the uncertainty process closer to A10. However, as the problem formulation to this point relies on a single recourse decision, these additional intermediate levels only serve to increase the dimensionality of the problem ($11^{n+1} + 1$ nodes where n is the number of intermediate waypoints considered) given the proposed decision process. The direct methodology introduces a separate consequence, however, as the distance between MWH and A10 is on the order of 1000 nm . Calculating the uncertainty process ξ in terms of a single stage transition significantly decreases the longitudinal resolution of the uncertainty model, as the uncertainty process would be assumed to hold a single (though still random) value between MWH and A10.

As one of the primary motivations for the uncertainty modeling process was to introduce region-specific information, this attribute of the direct calculation method is clearly undesirable. In order to utilize the longitudinal resolution of the uncertainty model while maintaining feasible dimensionality of the indirect approach, the number of first-stage scenarios must be sufficiently large to capture the evolution of the uncertainty process between MWH and A10 while being small enough to remain computational tractability. For certain classes of problems, namely convex problems with piecewise linear objectives and non-empty, bounded solution sets, the sample average approximation (SAA) [45] method may be used to reliably reduce the number of required scenarios [79, 54]. However, as previously discussed in Section 3.1.4, the RTA problem is ill-conditioned to a linearized formulation. As a result, the number of first stage scenarios will be parameterized as a variable of interest, and the effects of scenario size on the algorithm’s performance examined in a sensitivity analysis.

Assuming n first stage scenarios are required, the process of generating realizations is rather straightforward. The inverse transform method is used to sample the conditional cumulative mass functions generated using the initial distribution λ_{MWH} , one-step transition probability matrices, and Equation 12 at intermediate waypoints between MWH and A10 for both the East and North directional models. These samples form a series of uncertainty measurements along the route for each model direction $\xi_{A1}, \xi_{A2}, \dots, \xi_{A10}$, with the time history of samples constituting a single first-stage scenario $\xi_{[A10]}^t$. The sampling process is repeated until $t = n$ first stage scenarios have been generated. For each first stage scenario $\xi_{[A10]}^t$, the sampling process continues across $m \geq n$ second stage scenarios. A time history of uncertainty measurements $\xi_{A11}, \xi_{A12}, \dots, \xi_{BNA}$ is generated at intermediate waypoints between A10 and BNA; these samples constitute a single second stage scenario $\xi_{[BNA]}^\eta$, with the process repeating until $\eta = m$ second stage scenarios have been generated for each of the n first stage scenarios.

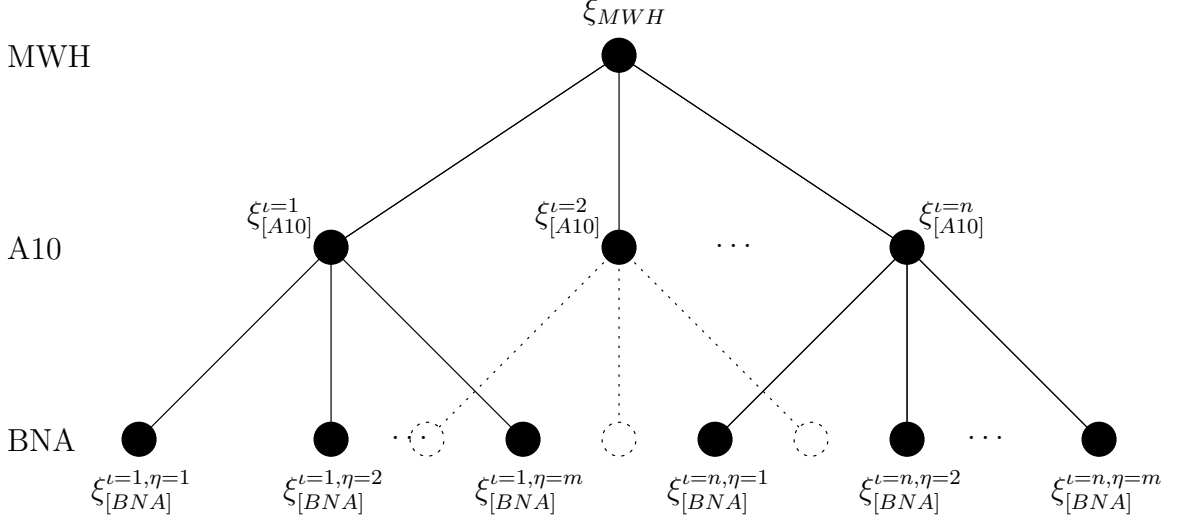


Figure 27: A two-stage scenario tree for a flight between KSEA and KATL.

The structure of the scenario tree with its $(n \times m + 1)$ nodes is illustrated in Figure 27, where each unique scenario is indexed by stage A10 node ι and stage BNA node η (note that branch $\xi_{[A10]}^{\iota=2}$ does exist, it was merely hidden to save figure space). Given this structure, a complete scenario is comprised of a series of sampled uncertainty values of the form $\xi_{[BNA]}^{\iota,\eta} = \{\xi_{MWH}, \xi_{A1}, \xi_{A2}, \dots, \xi_{A10}, \xi_{A11}, \xi_{A12}, \dots, \xi_{BNA}\}$

3.3.2 Constraints

Again consider the two-stage scenario posed in Figure 26. The task is to select a Mach number M_1 , that enables the aircraft to meet the assigned RTA at BNA given the aircraft is able to later change speeds at A10 in order to do so; this choice in initial Mach number M_1 should result in the minimum expected fuel burn for the entire flight given the wind forecast uncertainty scenarios possibly encountered along the route. This scenario can be constructed in terms of a multi-stage formulation as posed in Section 3.1.2:

$$\min_{M_1 \in \mathcal{M}_1} f_{MWH}(M_1) + \mathbb{E} [Q_{A10}(M_1, \xi_{[A10]})]$$

where the “cost-to-go” function can be expressed:

$$Q_{A10}(M_1, \xi_{[A10]}) = f_{A10}(M_2) + \mathbb{E} [Q_{BNA}(M_2, \xi_{[BNA]})]$$

Combining both stages, the problem follows:

$$\min_{M_1 \in \mathcal{M}_1, M_2 \in \mathcal{M}_2} f_{MWH}(M_1) + \mathbb{E} [f_{A10}(M_2) + \mathbb{E} [Q_{BNA}(M_2, \xi_{[BNA]})]] \quad (25)$$

where in this case the function f represents the fuel burned as a function of the chosen Mach number for the given flight segment ($f_{MWH}(M_1)$ would therefore represent the fuel burned along the segment from MWH to A10 given a Mach number of M_1). Furthermore, the expected value of the “cost-to-go” function $\mathbb{E} [Q(M, \xi)]$ represents the expected value of the fuel burned across the entire set of possible required speed changes at a given waypoint made necessary to correct for the effects of the forecast uncertainty process ξ on the estimated arrival time at the final waypoint required to satisfy the RTA.

The term $\mathbb{E} [Q_{BNA}(M_2, \xi_{[BNA]})]$ introduces an interesting issue when considering speed selection for RTA adherence. The problem could conceivably be refactored into an n -stage problem with as many intermediate RTA control points as desired. However, the n -th stage must enforce the aircraft arrive at the RTA waypoint with a given certainty or confidence. In practice, this will never be possible given the nature of the interaction between the forecast uncertainty process ξ and the RTA constraint. Consider the scenario presented in Figure 28 of an aircraft approaching an RTA control waypoint. As long as there is some distance between the aircraft’s current position and the control waypoint, there will be some arrival time uncertainty introduced by the wind forecast uncertainty process ξ . This arrival time uncertainty can be quantified in terms of a probability distribution describing the aircraft’s probable arrival times to the RTA fix rather than a singular estimated time of arrival. As the aircraft approaches the RTA control waypoint, the width of arrival time distribution will narrow as there is a shorter distance over which any forecast uncertainty may aggregate, with the distribution approaching the actual arrival time as the distance between the aircraft and the RTA control point approaches zero. Regardless of the number of problem stages, it will never be possible to ensure the aircraft arrives at the RTA




Figure 28: Effects of forecast uncertainty on arrival time distributions versus estimation distance.

waypoint at the assigned RTA given there is some measure of wind forecast uncertainty and distance remaining between the aircraft and the RTA control waypoint. As a result, the algorithm proposed in this work will assume that some level of air traffic controller or traditional control methodology intervention will be required at the tactical level as the aircraft approaches the neighborhood of the RTA fix in order to ensure the RTA is satisfied. The term $\mathbb{E} [Q_{BNA}(M_2, \xi_{[BNA]})]$ essentially quantifies the cost of this tactical level intervention, or n -th stage recourse decision.

It is not possible to model the cost of the n -th stage recourse decision. However, it is possible to reformulate the problem in an attempt to mitigate its effects. Again recall the generic formulation presented in Equation 25:

$$\min_{M_1 \in \mathcal{M}_1} f_{MWH}(M_1) + \mathbb{E} [f_{A10}(M_2) + \mathbb{E} [Q_{BNA}(M_2, \xi_{[BNA]})]]$$

To this point, no mention has been made of the constraints associated with the proposed model other than the implicit assumption that the aircraft must satisfy a specified RTA at the final waypoint, BNA. As the cost of the final recourse decision $\mathbb{E} [Q_{BNA}(M_2, \xi_{[BNA]})]$ is effectively unknown, a modification to the constraints is

made to mitigate the effects of the required action: rather than requiring that a set of speed decisions exactly satisfy the RTA, a set of decisions (M_1^ι, M_2) for $\iota \in 1, 2, \dots, n$ stage-one scenarios is considered to be a feasible policy $\mathbf{M}(\xi_{[BNA]}) \in \mathcal{M}_1 \times \mathcal{M}_2$ if $\left| \mathbb{E} \left\{ t_{BNA}(M_1^\iota, M_2, \xi_{[BNA]}^\iota) \right\} - t_{RTA} \right| \leq \epsilon$. In this instance, $t_{BNA}(M_1^\iota, M_2, \xi_{[BNA]}^{\iota, \eta})$ is the estimated arrival time of the aircraft to the RTA waypoint BNA given first stage Mach number M_1 indexed by stage one scenario ι , second stage Mach number M_2 , and a realization of the forecast uncertainty process $\xi^{\iota, \eta}$; $\left\{ t_{BNA}(M_1^\iota, M_2, \xi_{[BNA]}^\iota) \right\}$ is a sequence of arrival times of size $1 \times m$ indexed by stage two scenario η ; t_{RTA} is the required time of arrival at the RTA fix; and ϵ is a chosen tolerance. In other words, the recourse decision for each stage one scenario ι , M_2 , is chosen such that the expected value of the arrival time to the RTA fix due to the flight time estimation error effects introduced by ξ is sufficiently close to the RTA to some specified tolerance. The reader should again note that as a result of this modification, the choice in M_2 calculated by the RTA algorithm will not necessarily deliver the aircraft to the RTA fix exactly satisfying the RTA due to the effects of wind forecast uncertainty. Rather, it results in a choice in M_2 providing arrival times that on average satisfy corresponding RTA's over a series of operations along the same route; tactical level controller intervention may still be required on a per-flight basis.

Reformulating Equation 25 for completeness:

$$\begin{aligned}
& \min_{M_1 \in \mathcal{M}_1(\xi_{[A10]}), M_2 \in \mathcal{M}_2(\xi_{[A10]}, \xi_{[BNA]})} f_{MWH}(M_1) + \mathbb{E} [f_{A10}(M_2)] \\
& \text{s.t.} \quad \left| \mathbb{E} \left\{ t_{BNA}(M_1^\iota, M_2, \xi_{[BNA]}^\iota) \right\} - t_{RTA} \right| \leq \epsilon \quad (26) \\
& \quad \quad \quad \forall \iota \in 1, 2, \dots, n
\end{aligned}$$

Given the formulation posed in Equation 26, focus must now shift to the evaluation of the formulation objective and constraint parameters for fuel burn and flight time.




Figure 29: Wind triangle.

3.3.3 Aircraft Model

Numerical flight time estimation in the context of the RTA algorithm developed in this work is a rather straightforward process. Given a fixed ground track as constructed in Section 2.3.3.2, estimating an along-track flight time is simply a matter of mapping the aircraft's true airspeed to its groundspeed given the current wind forecast and wind uncertainty scenario. Recall that the aircraft's groundspeed vector is a function of its true airspeed vector as well as the incident wind vector at its current position:

$$\mathbf{v}_g = \mathbf{v}_{TAS} + \mathbf{v}_{wind},$$

where the wind vector in this case is the sum of a nominal wind forecast and some forecast uncertainty process:

$$\mathbf{v}_{wind} = \mathbf{v}_{forecast} + \boldsymbol{\xi}$$

As the aircraft is required to maintain its lateral track along its planned route, the aircraft's groundspeed vector should track the planned lateral path. Consequently, a heading offset λ is required to correct for drift induced by the incident wind vector (this scenario is illustrated in Figure 29). The incident wind vector is decomposed into headwind and crosswind components using the angle difference β between the incident wind speed vector and the desired ground track heading θ :

$$\mathbf{W}_{head} = |\mathbf{v}_{wind}| \cos \beta$$

$$\mathbf{W}_{cross} = |\mathbf{v}_{wind}| \sin \beta$$

and the resultant groundspeed follows:

$$v_g = |\mathbf{v}_{TAS} - \mathbf{W}_{cross}| - |\mathbf{W}_{head}| \quad (27)$$

A further modification is required, however, in order to model the known behavior of current flight management systems as it pertains to the handling of forecast winds. Current systems blend known winds at the aircraft's present position with forecast winds at the next waypoint using simple linear interpolation. In order to mimic this wind mix behavior, the wind speed at any given waypoint is calculated as the average value between the known wind at the aircraft's current waypoint i , and the forecasted winds at the next waypoint $i + 1$:

$$\mathbf{v}_{wind,i} = \frac{1}{2} \cdot (\mathbf{v}_{actual,i} + \mathbf{v}_{wind,i+1})$$

where from Equation 5:

$$\mathbf{v}_{actual,i} = \mathbf{v}_{forecast,i} + \boldsymbol{\xi}_i^t$$

and $\boldsymbol{\xi}^t$ is the current forecast uncertainty scenario.

As the groundspeed can now be calculated at any waypoint along the planned route given the planned heading, the aircraft's true airspeed vector, a wind forecast, and a forecast uncertainty scenario, the flight time along a given segment distance can be numerically estimated based on a discretization of the flight path into smaller segments. The flight time estimation routine used by the RTA algorithm first discretizes the complete flight path into the same 100 *nm* segments utilized by the forecast uncertainty model developed in Chapter 2. The routine then calculates the flight time for each segment, finally summing the segment flight times to find the total path flight time. This process is summarized in Algorithm 2.

Estimating aircraft fuel expenditure is performed via a first principles model of aircraft performance. The lift required to maintain steady-level flight is equal to the instantaneous weight of the aircraft W :

$$L_{req} = W = \frac{1}{2} \rho_\infty \cdot |\mathbf{v}_{TAS}|^2 \cdot C_{L,req} \cdot S \quad (28)$$

Algorithm 2 Estimates flight time along a path given discretized path segments.

given $v_{forecast}$ {wind forecast by route segment}
given ξ_i^t {wind forecast uncertainty scenario by route segment}
given \mathbf{v}_{TAS} {true airspeed vector}
given θ {desired ground track heading by route segment}
given d {route segment lengths}
init $T \leftarrow 0$ {total path flight time}
for all flight path segments in set $\{1 \dots i \dots (N - 1)\}$ **do**
 $\mathbf{v}_{actual,i} \leftarrow \mathbf{v}_{forecast,i} + \boldsymbol{\xi}_i^t$
 $\mathbf{v}_{wind,i} \leftarrow \frac{1}{2} \cdot (\mathbf{v}_{actual,i} + \mathbf{v}_{wind,i+1})$
 $\beta \leftarrow \angle(\mathbf{v}_{wind}) - \theta$
 $\mathbf{W}_{head} \leftarrow |\mathbf{v}_{wind,i}| \cos \beta$
 $\mathbf{W}_{cross} \leftarrow |\mathbf{v}_{wind,i}| \sin \beta$
 $v_{g,i} = |\mathbf{v}_{TAS} - \mathbf{W}_{cross}| - |\mathbf{W}_{head}|$
 $t_i \leftarrow v_{g,i}/d_i$ {segment flight time}
 $T \leftarrow T + t_i$
end for
return T

where ρ_∞ is the atmospheric density of the free stream air, $C_{L,req}$ is the required coefficient of lift to maintain steady-level flight, and S is the planform area of the aircraft's wings. Given the weight of the aircraft is known, the required coefficient of lift follows:

$$C_{L,req} = \frac{2 \cdot W}{\rho_\infty \cdot |\mathbf{v}_{TAS}|^2 \cdot S} \quad (29)$$

Using a proprietary drag polar [Boeing Performance Engineer's Manual: 737-300/CFM56-3B-2], the required coefficient of lift $C_{L,req}$ is mapped to the corresponding coefficient of drag C_D as a function of both the required coefficient lift and the aircraft's Mach number M . As the instantaneous thrust must equal the total drag in order to maintain steady-level flight, the required thrust T_{req} is found:

$$T_{req} = D = \frac{1}{2} \rho_\infty \cdot |\mathbf{v}_{TAS}|^2 \cdot C_D \cdot S \quad (30)$$

Proprietary fuel burn rate tables [Boeing Performance Engineer's Manual: 737-300/CFM56-3B-2] are then used to map the required thrust T_{req} to the required fuel burn rate \dot{m} given the aircraft's cruising altitude h and the aircraft's Mach number.

Algorithm 3 Estimates fuel expenditure given initial aircraft weight and travel time.

```
given  $h$  {wing planform area}
given  $W_0$  {initial cruising weight}
given  $T$  {total travel time}
given  $\delta t$  {time step}
given  $h$  {altitude}
load  $\rho_\infty \leftarrow ICAO(h)$  {air density at flight level  $h$  [41]}
load  $a_\infty \leftarrow ICAO(h)$  {speed of sound at flight level  $h$  [41]}
init  $t \leftarrow 0$ 
init  $W \leftarrow W_0$ 
 $M \leftarrow |\mathbf{v}_{TAS}|/a_\infty$ 
while  $t < T$  do
   $C_{L,req} \leftarrow (2 \cdot W)/(\rho_\infty \cdot |\mathbf{v}_{TAS}|^2 \cdot S)$ 
  load  $C_D(M, C_{L,req})$  {mapped from drag polar}
   $T_{req} \leftarrow \frac{1}{2} \rho_\infty \cdot |\mathbf{v}_{TAS}|^2 \cdot C_D(M, C_{L,req}) \cdot S$ 
  load  $\dot{m}(h, M, T_{req})$  {mapped from fuel burn rate tables}
   $W \leftarrow W - \dot{m}(h, M, T_{req}) \cdot \delta t$ 
   $t \leftarrow t + \delta t$ 
end while
return  $(W_0 - W)$ 
```

The model developed in Equations 28 - 30 correspond to instantaneous values for the aircraft's weight. In order to quantify the impact of varying aircraft weight as fuel is burned along-route, the total fuel expenditure must be calculated iteratively based on a discretization of the aircraft's total travel time T . Algorithm 3 summarizes this process assuming the aircraft's initial cruising weight W_0 , true airspeed, and total travel time T are known. Atmospheric parameters are defined as a function of the ICAO standard atmosphere and the cruising altitude of the aircraft [41].

The availability of propriety aircraft performance data makes fuel burn estimation significantly simpler. The RTA algorithm implementation presented in this research effort makes use of performance data for the Boeing 737-300/CFM56-3B-2, though similar tables are available for a variety of Airbus and Boeing aircraft. In the absence of such data, a more robust performance model such as EUROCONTROL's Base of Aircraft Data (BADA) [26] would be required in order to generate sufficiently accurate estimates of the aircrafts fuel burn characteristics. However, as the RTA algorithm is

designed for an FMS-based implementation rather than a ground system, assuming access to relevant vehicle performance data is reasonable.

3.3.4 Search Heuristic

The formulation posed in Equation 26 allows for the problem to be solved in a series of unique subproblems according to stage: the second stage strictly enforces the expected arrival time constraint whereas the first stage (in consideration of nested second stage solutions) deals entirely with Mach selection for fuel burn minimization. The proposed RTA algorithm leverages this staged problem structure along with aircraft performance characteristics to efficiently search forecast uncertainty scenario sets generated using the methods outlined in Section 3.3.1 for a fuel optimal implementable policy. The algorithm architecture is depicted in Figures 30-32. The outermost level is described in Figure 30; this level includes the primary Mach number search. The objective function (fuel burn) calculation is similarly shown in Figure 30. Lastly, the recourse decision search based on the constraint described in Equation 26 is outlined in Figure 32. Figures 30 and 32 highlight search functions which will be further detailed in Section 3.3.4.1.

Beginning with the outer most level depicted in Figure 30, the RTA algorithm begins by calculating a feasible solution. For simplicity, this initial Mach number is chosen such that it satisfies the RTA given perfect forecast information (i.e.: $\mathbf{w}_{forecast} = \mathbf{w}_{actual}$, or $\boldsymbol{\xi} = 0$). The feasible Mach solution is next used to calculate the expected fuel burn across all $n \times m$, stage one and stage two scenarios. The reader should note from Figure 30 that this process involves both the objective calculation subroutine described in Figure 31 and the RTA adherence subroutine described in Figure 32.

To calculate the fuel burn for the feasible Mach solution, first the flight time from MWH to A10 for each realization of the forecast uncertainty model $\xi_{[A10]}^\ell$, $\ell \in$

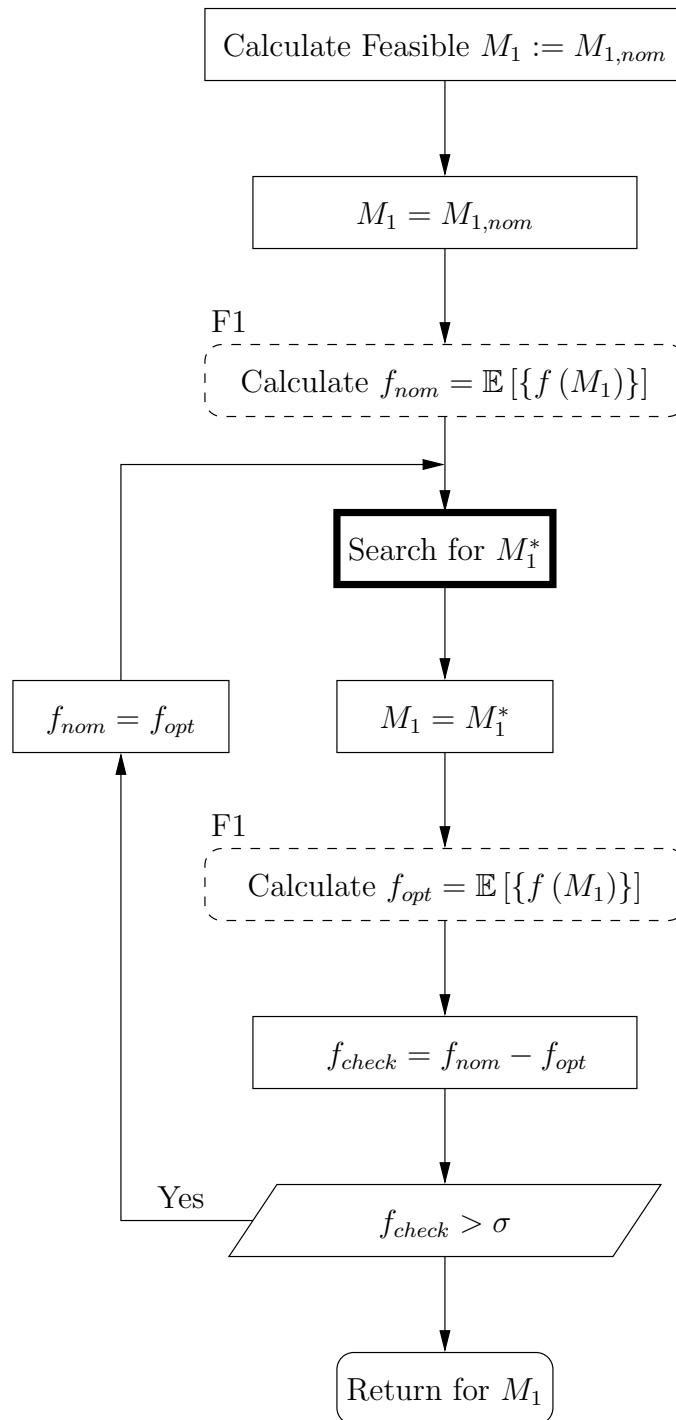


Figure 30: Stage one fuel burn minimization subproblem functional diagram.

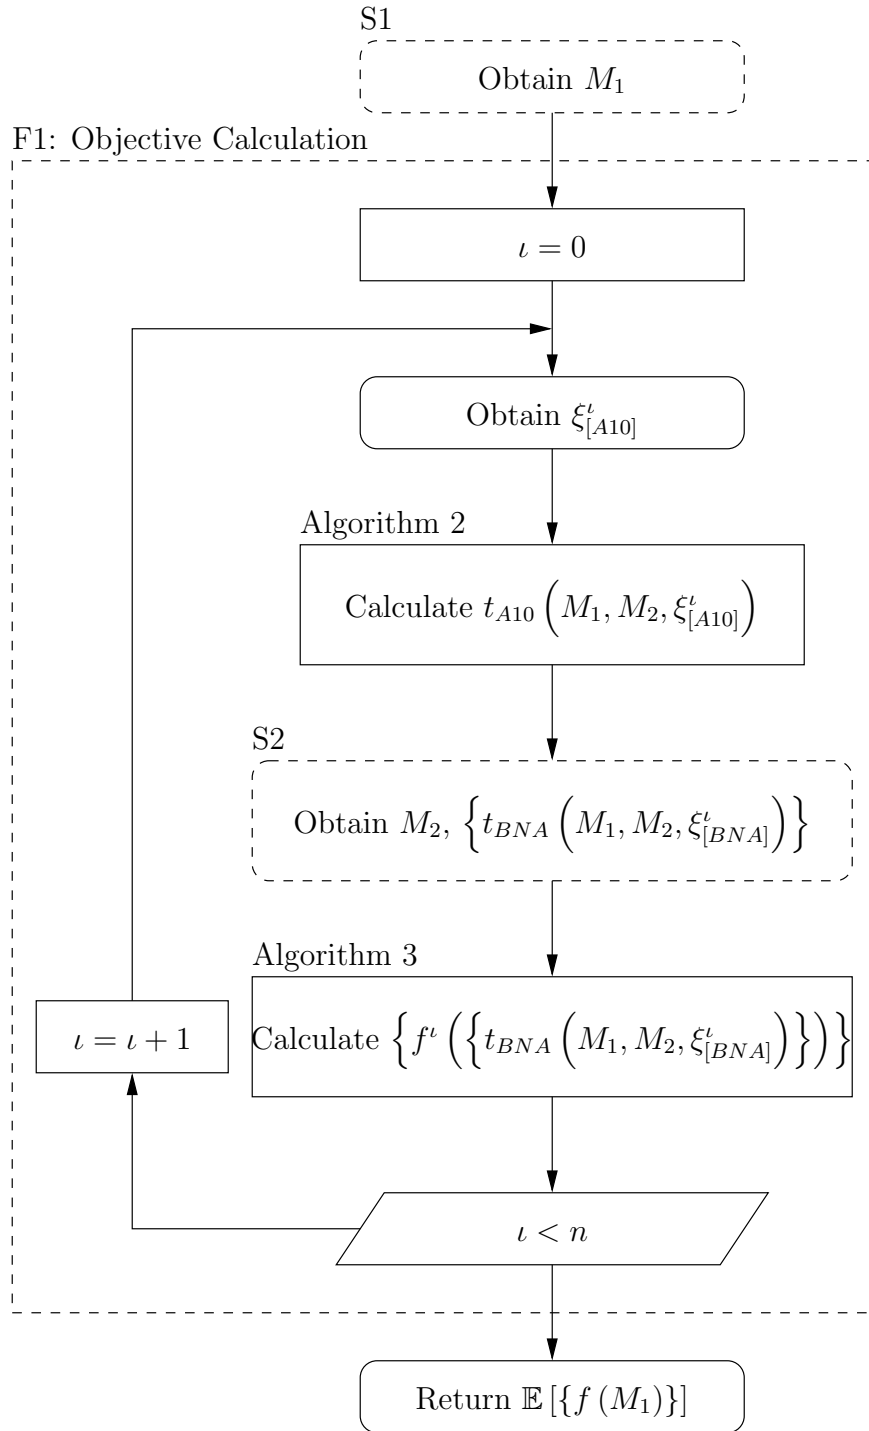


Figure 31: Stage one objective calculation functional diagram.

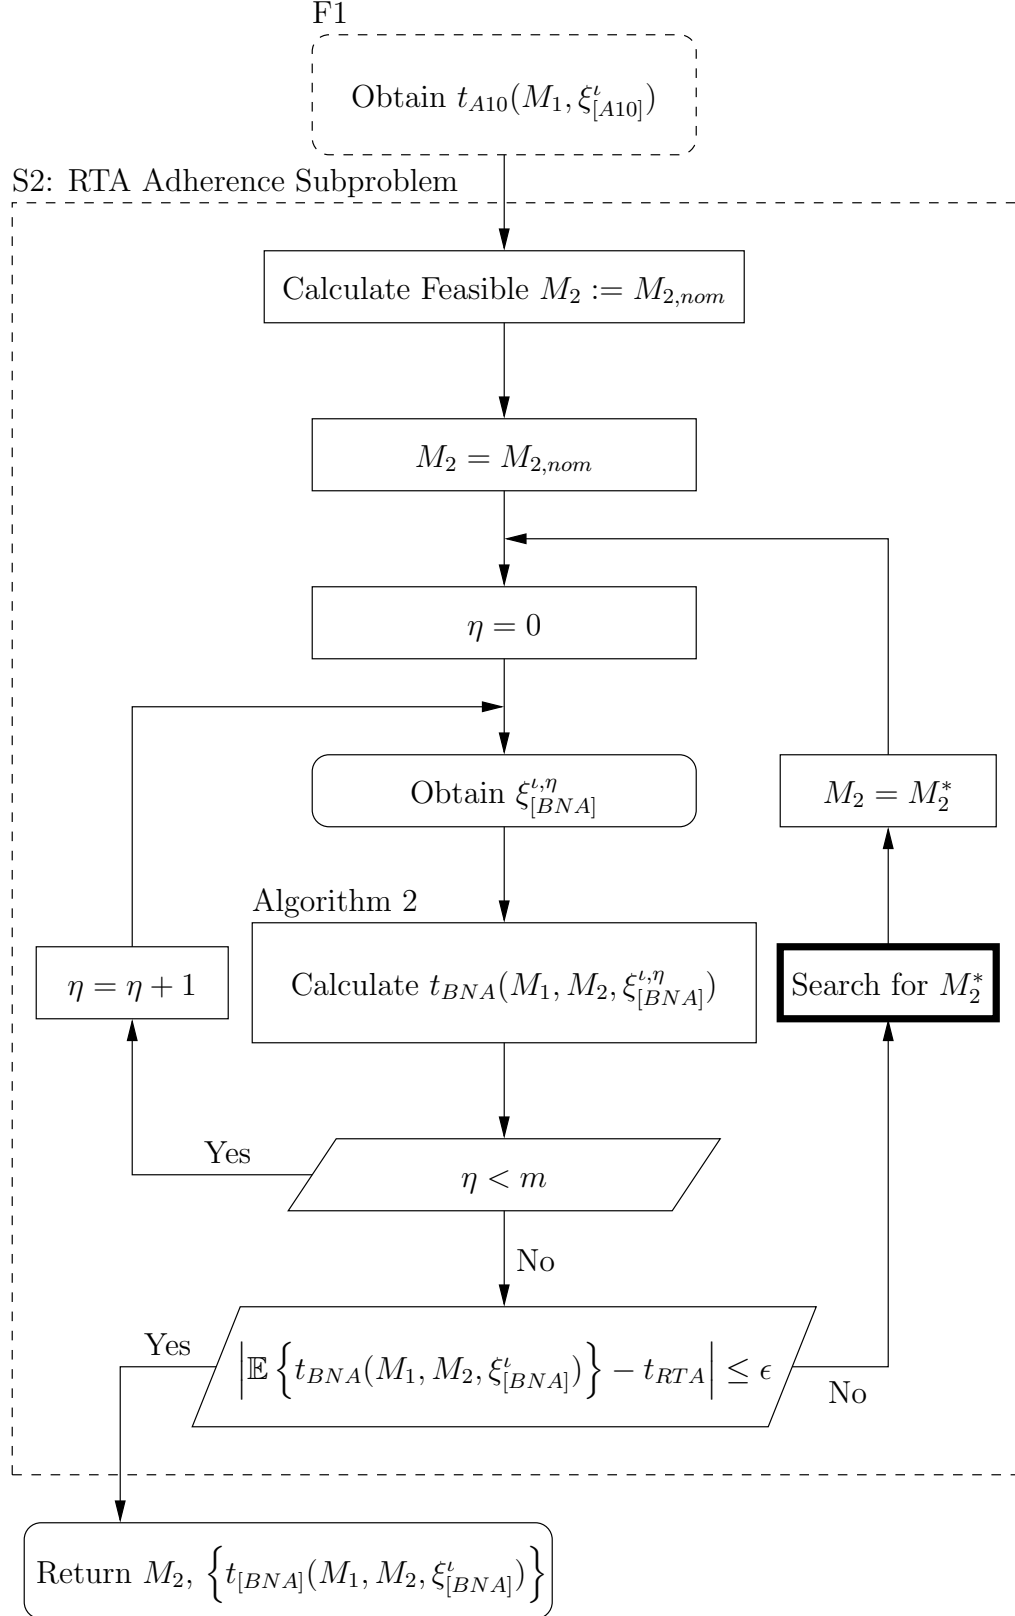


Figure 32: Stage-two RTA adherence subproblem functional diagram.

$1, 2, \dots, n$, is calculated using the method outlined in Algorithm 2 and depicted in Figure 31. The result is a sequence of n arrival times to A10, $\{t_{A10}(M_1, \xi_{[A10]})\}$. Each arrival time represents a possible stage one outcome based on the initial feasible Mach solution and the impact of the forecast uncertainty process between MWH and A10, $\xi_{[A10]}$. The RTA constraint defined in Equation 26 must now be applied to each of the stage one outcomes by means of a speed change at A10; this process is performed by the RTA adherence subroutine labeled S2, which is further described in Figure 32.

Calculating the second stage speed begins in the same manner as the outermost layer of the algorithm. For each of the stage one outcomes indexed by stage one scenario ι , a feasible Mach solution is calculated such that given perfect forecast information, the aircraft satisfies the RTA at BNA given its initial estimated arrival time to A10, $t_{A10}(M_1, \xi_{[A10]}^\iota)$. The stage two feasible Mach is then used to estimate the time of arrival at BNA again using Algorithm 2 for each second stage realization of the forecast uncertainty process $\xi_{[BNA]}^{\iota, \eta}$, $\eta \in 1, 2, \dots, m$. It should be clear to this point that there now exist $n \times m$ estimated arrival times to BNA, indexed first by the first stage scenario index ι , then by the second stage scenario index η . The constraint developed in Section 3.3.2 holds that for each subset of the arrival times indexed by a single stage one scenario, the expected value of the arrival time must fall sufficiently close to the RTA. Enforcing this constraint therefore requires calculating a second stage Mach, M_2 , that shifts the expected value of each subset of arrival times to the RTA at BNA *for each of the stage one scenarios*. Determining this Mach number, defined as M_2^* in Figure 32 is the focus of Section 3.3.4.1. For now, assume that once this value is determined, the final step of the RTA adherence subroutine is to return the second stage Mach number M_2 satisfying the RTA constraint as well as the m estimated arrival times to BNA.

Now that a value for M_2 has been found for each of the n stage one scenarios that satisfies the RTA constraint, and arrival times to BNA calculated based on

the stage one feasible Mach solution and the RTA-required stage two Mach number, Algorithm 3 is used to calculate the fuel burn for all $n \times m$ scenarios. The expected fuel burn across the full range of scenarios is then returned. This value for expected fuel burn represents the average fuel burn across all generated forecast uncertainty scenarios given the aircraft flies from MWH to A10 at the initial feasible stage one Mach number, then updates its speed at A10 in order to enforce the RTA constraint at BNA. Furthermore, the corresponding solution pair $\mathbf{M} = (M_1, M_2)$ is a feasible policy, though not necessarily an optimal policy. As one can see in Figure 30, a search is next performed on Mach number by a “black box” search routine to improve the initial feasible solution used by the algorithm, the process repeated, and the expected fuel burn results given the updated Mach number compared to the prior iteration. Once the difference in expected fuel burn between iterations falls below a specified tolerance (defined as σ), the algorithm is considered to have converged on a value for the first stage Mach number, M_1 , and the resultant policy \mathbf{M} considered to be an optimal policy.

3.3.4.1 Mach Searches

To this point, the algorithm as depicted in Figures 30-32 strictly lays out a framework for enumerating the forecast uncertainty model scenarios in a structured manner while enforcing the RTA constraint given an initial feasible first stage Mach solution, M_1 . The real power of the algorithm lies in the two Mach searches performed by the outermost layer of the algorithm described in Figure 30 and the inner most layer described in Figure 32.

First, consider the Mach search referenced in Figure 30. The goal of the routine is to generate a stage one Mach number that results in an improved value for the expected fuel burn across all $n \times m$ scenarios considering the possible recourse actions required at A10 when compared to its input, an initial feasible value for M_1 . Recall

from section 3.1.4 that fuel burn rate is approximately quadratic with Mach number. Intuitively, one would assume the expected fuel burn value returned by the algorithm objective subroutine to be quadratic with Mach number as well. Accordingly, the stage one search routine simply parameterizes the expected fuel burn as a quadratic function of the stage one Mach number. The input M_1 value is first perturbed in both directions, resulting in two additional values for M_1 in the neighborhood of the original M_1 value. The expected fuel burn for each perturbed Mach value is then calculated using the objective calculation subroutine described in Figure 31. A three-point quadratic approximation of expected fuel burn versus Mach number is then constructed, and the closed-form minimum of the function evaluated:

$$M^* = \frac{f_\ell + f}{2} - \frac{s_\ell \cdot (f_h - f_\ell)}{2 \cdot (s_h - s_\ell)} \quad (31)$$

with:

$$\begin{aligned} f_\ell &= \mathbb{E}[\{f(M_1 - \delta)\}] \\ f &= \mathbb{E}[\{f(M_1)\}] \\ f_h &= \mathbb{E}[\{f(M_1 + \delta)\}] \\ s_h &= (f_h - f) / ((M_1 + \delta) - M_1) \\ s_\ell &= (f - f_\ell) / (M_1 - (M_1 - \delta)) \end{aligned}$$

where M_1 is the input Mach number, δ is a small perturbation on the order of $M = 0.05$, and M^* is the minimum of the quadratic approximation. This minimum Mach value is then set as the new initial feasible solution, and the process repeated until the expected fuel burn converges.

The second Mach search, referenced in Figure 32, performs in nearly the same manner. However, in this case the distance between the mean arrival time of the stage two scenarios and the RTA is parameterized and fit to a quadratic approximation as

a function of the second stage Mach number:

$$\begin{aligned}
 f_\ell &= \left| \mathbb{E} \left\{ t_{BNA}(M_1, M_2 - \delta, \xi_{[BNA]}^\ell) \right\} - t_{RTA} \right| \\
 f &= \left| \mathbb{E} \left\{ t_{BNA}(M_1, M_2, \xi_{[BNA]}^\ell) \right\} - t_{RTA} \right| \\
 f_h &= \left| \mathbb{E} \left\{ t_{BNA}(M_1, M_2 + \delta, \xi_{[BNA]}^\ell) \right\} - t_{RTA} \right| \\
 s_h &= (f_h - f) / ((M_2 + \delta) - M_2) \\
 s_\ell &= (f - f_\ell) / (M_2 - (M_2 - \delta))
 \end{aligned}$$

where Equation 31 still holds for the minimum of the approximation. Again, this routine accepts an initial feasible solution for M_2 , and the resulting value for M_2 is meant to minimize the difference between the mean arrival time to the RTA waypoint, BNA, and the RTA; this process iterates until it fails to improve the time difference. As an additional note to the reader, a quadratic approximation for the second stage is not necessarily the only acceptable searching routine to use for this purpose. Whereas a quadratic approximation leverages the quadratic approximation between Mach number and fuel burn rate for the stage one Mach search, a quadratic fit routine was selected for this purpose as the prior search routine had already been implemented.

Table 6: Notional RTA scenarios along sample routes.

Route	Origin	Destination	Flight Level	Decision	RTA
1	KSEA	KATL	370	A09	03:45:00 GMT
2	KATL	KSEA	360	A09	03:55:00 GMT
3	KLAX	KATL	370	A08	03:00:00 GMT
4	KATL	KLAX	360	A08	03:45:00 GMT
5	KEWR	KATL	360	A03	01:30:00 GMT
6	KATL	KEWR	350	A03	00:58:30 GMT

3.4 *Sample Results*

As any result produced by the RTA algorithm is the product of a specific scenario construction, results are presented for notional RTA scenarios along the sample routes provided in Appendix A relative to a specific series of RUC-20 forecast releases. Each of the notional RTA scenarios hold the following characteristics in common:

1. The aircraft reaches its initial cruise waypoint at 2009-08-31 00:00:00 GMT.
2. The aircraft reaches its initial cruise waypoint at 70% of its maximum gross takeoff weight.
3. The forecast uncertainty model is initialized to represent a zero-error initial distribution (i.e.: the difference between the winds measured at the first cruise waypoint and the wind forecast lies in the range $[-1, 1)$).
4. The forecast uncertainty model used represents the “base” scenario discussed in Chapter 2.
5. The recourse decision occurs at a point roughly two-thirds along the length of the route.
6. The algorithm considers $n = 100$ stage-one and $m = 100$ stage two forecast uncertainty scenarios, where $|\xi| = 10,000$.

Table 6 lists a series of such scenarios along the sample route, where the "Decision" column lists the waypoint separating the problem stages (the reader should be aware by now that the recourse speed decision is made at this waypoint). Table 7 lists the corresponding algorithm solutions, where $M_{1,nom}$ is the required cruise speed to satisfy the RTA given no forecast uncertainty and M_1^* is the calculated optimal stage one cruise speed. The speed advisories listed in Table 7 are of diminished value, however, given no baseline metric of fuel burn improvement. In order to generate baseline values for comparison, a "traditional" dead-band RTA logic was also implemented to demonstrate the benefit of the fuel planning algorithm. An overview of this algorithm is provided in Algorithm 4; any "calculate" subroutines have already been developed in Section 3.3.3. As a reminder of the traditional approach to RTA functionality as introduced in Chapter 1, the estimated time of arrival (ETA) to the final waypoint is monitored as an aircraft progresses along its route. At any point, if the projected ETA diverges from the RTA to a degree larger than a specified tolerance, the cruising Mach number is adjusted. Using the logic developed in Algorithm 4, three more columns are generated for Table 7: $\mathbb{E}[f_{M_1^*}]$ is the expected fuel burn across all scenarios given a calculated optimal stage one Mach number found by the stochastic algorithm, $\mathbb{E}[f_M]$ is the expected fuel burn across all scenarios given a traditional RTA approach as detailed in Algorithm 4 (note that \mathbf{M} in this case is a vector quantity as there are multiple possible cruising Mach numbers along a given scenario), and δ_f is the difference between $\mathbb{E}[f_{M_1^*}]$ and $\mathbb{E}[f_M]$ (a negative value represents an expected savings in favor of the stochastic algorithm in terms of pounds of fuel).

In examining Table 7, the reader should reference the scenario conditions detailed at the start of this section. Recall that these sample scenarios were built using a wind forecast uncertainty model initialized with a zero-error initial distribution, meaning that forecasted winds along the route added very little uncertainty to the solution

Algorithm 4 Dead-band speed controller for RTA flight mode.

```

given  $\epsilon$  {control gain}
given  $T_0$  {current clock time}
given  $RTA$  {required time of arrival at final route waypoint}
calculate  $M_0(T_0)$  {cruise Mach required to satisfy the RTA from time  $T_0$ }
init  $T \leftarrow T_0$  {flight clock}
init  $M \leftarrow M_0$  {cruise Mach number}
for all flight path segments in set  $\{1 \dots i \dots (N - 1)\}$  do
  calculate  $t_i(M)$  {flight time to segment end given cruise Mach  $M$ }
   $T \leftarrow T + t_i$ 
  calculate  $ETA_i(M, T)$  {estimated time of arrival to final route waypoint}
  if  $|ETA_i(M, T) - RTA| > \epsilon$  then
    calculate  $M_i(T)$  {Mach required to satisfy the RTA from segment  $i$ }
    init  $M \leftarrow M_i$ 
  end if
end for

```

Table 7: Notional RTA scenario results.

Route	$M_{1,nom}$	M_1^*	$\mathbb{E} [f_{M_1^*}]$ (lbs)	$\mathbb{E} [f_M]$ (lbs)	δ_f (lbs)
1	0.715	0.735	20,515	20,584	-69
2	0.710	0.730	21,164	21,238	-74
3	0.715	0.730	16,586	16,618	-32
4	0.715	0.735	20,268	20,350	-82
5	0.720	0.730	8,053	8,071	-18
6	0.710	0.715	5,452	5,458	-4

space. This fact highlights a characteristic of the standard RTA approach previously introduced in Chapter 1: traditional dead-band control solutions force aircraft to operate outside of their optimal performance envelopes [82]. Though the forecast uncertainty model had little impact on the scenarios, there were still significant fuel savings to be found simply by not constraining the aircraft to continuously track a suboptimal cruising speed. Alternatively, the author’s approach allows the FMS to consider the cost of cruising at a slightly better Mach number for a portion of the route given a corrective speed choice can be made later to satisfy the RTA; furthermore, the fuel cost of this correction is considered in the initial Mach selection.

Routing plays an interesting role in the performance of the algorithm. An obvious result is that more fuel is saved along longer flight paths, as small speed changes have more time to affect change in the amount of fuel burned along the route. One may initially consider the smaller fuel burn difference along route 6 between KATL and KEWR to be an insignificant result in comparison to the results for the transcontinental scenarios. Further inspection reveals however, that the small fuel savings provided by the stochastic algorithm along route 6 is actually a statistically significant result. The standard error of the difference in mean fuel burns along route 6 was found to be 0.306 *lbs*, corresponding to a 99% confidence interval for the difference in the mean fuel burns of 4 ± 0.860 *lbs*.

A more nuanced result was introduced in Section 2.4.1. Notice that in general, fuel savings are greater along routes opposing the jet stream’s natural direction (routes 4 and 5 in particular). Now consider that an increase in the stage one Mach number from the nominal value allows the aircraft to operate at a better Mach number from a fuel burn rate standpoint, but requires a larger decrease in speed at the recourse waypoint in order to satisfy the RTA. Aircraft will tend to experience a stronger headwind or weaker tailwind along routes opposing the jet stream, meaning the magnitude of this speed decrease is slightly reduced. As a result, the aircraft continues

to operate more closely to its fuel-preferred Mach number. This effect is particularly noticeable along routes 3 and 4. As the routing is nearly direct in both directions between KATL and KLAX, avoiding a headwind is virtually unavoidable along route 4 (as reflected in the increased flight time and fuel burn); this attribute of the route is exploited by the algorithm, however, to find significantly more fuel savings when compared to the converse routing.

3.4.1 Time to Solution

Questions regarding the complexity of linear programming algorithms have been raised since the 1950's. Klee and Minty demonstrated that the simplex algorithm could at worst be expected to perform in exponential time in 1970 [47]; Khachiyan demonstrated that linear programming problems can be solved in polynomial time relative to the length of the input data using ellipsoid methods in 1979 [46]; and Karmarkar developed a polynomial time interior point algorithm for linear programs in 1984 [43]. Linear programming methods are used as a comparison for the author's search heuristic as the alternative methods explored in Sections 3.1.1-3.1.4 utilized such techniques. However, it is important to note that the algorithm developed in this research effort is strictly a heuristic, meaning optimal solutions found by the algorithm have not been proven to be globally optimal as in the case of the linear methods discussed prior.

As detailed in Section 3.3.4.1, the heuristic makes use of a three point quadratic approximation to search for the fuel-minimum stage one Mach number. Given fuel burn rate is roughly quadratic with Mach number, one should expect the quadratic approximation of a quasi-quadratic function to provide the minimum solution in a single iteration. Absent numerical instabilities introduced by the aircraft performance data, this is often the case. Intrinsicly, one would therefore expect the algorithm to reach a solution in linear time, as the solution time is simply bound by the time

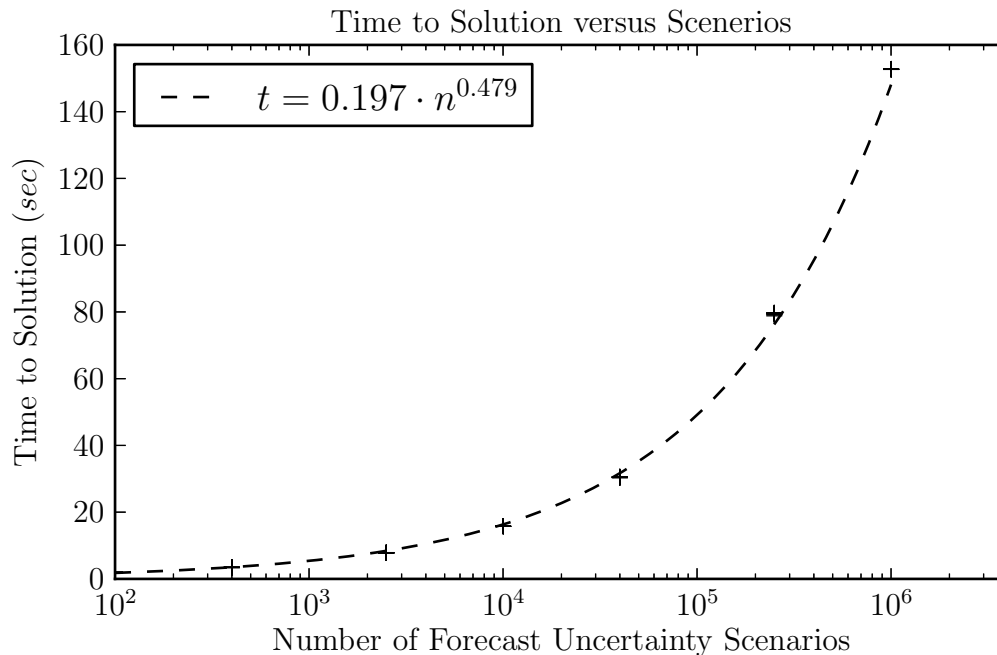


Figure 33: Algorithm time to solution as a function of scenario set size for sample scenario 1.

it takes to enumerate the tree completely three times. However, if one considers the structure of the scenario tree, a different conclusion is reached. Recall from Section 3.3.4 that each first stage branch is enumerated independently and the flight time from the first waypoint to the decision waypoint given the first stage wind forecast uncertainty scenario calculated. Further flight time calculations are then performed on each of the stage two forecast scenarios for the individual stage one branches. However, the stage two calculations only consider the flight path between the decision waypoint and the final RTA waypoint. Accordingly, the forecast uncertainty scenario tree is not enumerated completely as a function of individual scenarios, rather as a function of the individual stage one branches. The algorithm time to solution is therefore driven by the number of stage one scenarios rather than the total number of scenarios. Moreover, the algorithm provides a solution in better than linear time. An examination of the algorithm run time is provided in Figure 33, which depicts the

Table 8: Sensitivity to decision waypoint location.

Route	Decision Waypoint	M_1^*	$\mathbb{E} [f_{M_1^*}]$ (lbs)
1	A06	0.750	20,508
	A09	0.735	20,515
	A12	0.735	20,521
2	A06	0.750	21,149
	A09	0.730	21,164
	A12	0.720	21,178
3	A05	0.740	16,570
	A08	0.730	16,586
	A11	0.730	16,598
4	A05	0.745	20,251
	A08	0.735	20,268
	A11	0.735	20,333

times to solution for 10 runs of sample scenario 1 (defined in Table 6) as a function of the size of the forecast uncertainty scenario set.² As one can see in Figure 33, there is little variation in run time for problems of the same size. Furthermore, the algorithm provides a solution in fractional power time. Though a slight improvement over linear run time, fractional power time represents a significant run time improvement when compared to a polynomial time solution.

3.4.2 Sensitivity to Decision Waypoint

Table 8 lists a series of sensitivity studies examining the effects of decision point placement on the calculated solution for sample scenarios 1-4 as constructed in Table 6 (routes between KATL and KEWR were omitted as the route length is too short to produce any perceptible change in result). In each case, as the stage one decision

²The particular algorithm implementation was developed using a mixture of Python 2.7 (with Numpy 1.7) and standard C. Though Python introduces some computational overhead as an interpreted language, this choice was made to both ease flexibility of development and respect the technology debt inherited by the weather modeling portion of this research effort. Best-practices regarding minimizing interpreter overhead were respected. However, one should expect significant solution time improvement given a purely C (or in the case of an FMS implementation, Ada) implementation. The algorithm was run on a Core i5 processor with a clock speed of 1.8GHz in a single processing thread; this implementation was purposefully not parallelized to respect the nature of on-board systems.

length is reduced, the recommended stage one Mach number is increased. Again, recall that aircraft Mach number is roughly quadratic with Mach number and as a result, there exists a Mach number at which the aircraft’s fuel burn rate achieves a minimum value. Given the 70% weight fraction defined in scenario definition at the beginning of Section 3.4, this minimum Mach lies near Mach 0.77 (this value was verified, however the author cannot disclose a plot of fuel burn versus Mach number as this information constitutes protected intellectual property). The trend in increasing Mach is a direct result of the aircraft’s optimal operating Mach lying at 0.77. Given the aircraft has a longer stage two distance, a less dramatic recourse decision is required to counter the time effects of the increased stage one Mach number.

Further consideration must be given to decision waypoint placement, however, when one considers the algorithm formulation developed in Section 3.3.2. Recall that the algorithm considers the RTA constraint in terms of expected value as there will always be some measure of time uncertainty between the aircraft’s estimated arrival time to the RTA waypoint and the RTA so long as some distance remains between the two. Accordingly, one can consider the placement of the decision waypoint as an analogue for acceptable RTA adherence risk. The further the decision waypoint lies from the RTA waypoint, the higher risk of the aircraft arriving to the RTA waypoint within some small neighborhood of the RTA. However, moving the decision waypoint further from the RTA waypoint may allow for fuel saving opportunities.

3.4.3 Sensitivity to Forecast Uncertainty Conditions

The algorithm demonstrated very little deviation in performance given a variety of initial distributions used to initialize the forecast uncertainty model. As a specific example, distributions representing initial conditions of error magnitude proportional to 7.7 *m/s* (the level at which forecast error becomes detrimental to air traffic management operations—see Section 2.1.1.1), and the algorithm produced nearly identical

results to those presented in Table 7. However, this result is a shortcoming of the forecast uncertainty model rather than the algorithm and will be discussed in further detail in Chapter 4.

CHAPTER IV

CONCLUSIONS

The shortcomings of current flight management system (FMS) implementations of the required time of arrival (RTA) mode in the presence of wind forecast uncertainty were introduced and discussed in Chapter 1. The clear need for an on-board system capable of making fuel-optimal speed decisions to enable the controlled time arrivals in the presence of forecast uncertainty was identified and a two-fold solution proposed. First, a location-specific, data-driven forecast uncertainty model was developed in order to capture regional uncertainty characteristics for the Rapid Update Cycle (RUC20) forecast uncertainty product. Constructed as an inhomogeneous Markov model, the forecast uncertainty model served as a scenario generation mechanism for a two-stage stochastic algorithm. The stochastic algorithm aimed to calculate the stage one cruise speed that ensured the aircraft's expected arrival time to the RTA waypoint was within a certain tolerance of the RTA while providing the minimum expected fuel burn along the entire route.

An examination of the literature illuminated two primary gaps in current approaches to forecast uncertainty modeling: all approaches to date were unable to resolve error introduced by the characteristics of the sensor network (i.e.: errors of scale or representativeness), and no other efforts had genuinely attempted to model the regional characteristics of the RUC20 model, with most works relying on statistical information generated during a year-long evaluation of the RUC product in the Denver center airspace. As the author had already identified regional variations in the RUC20 product in prior research efforts, a model capturing regional forecast uncertainty characteristics rather than relying on statistics from a single region was

deemed necessary, especially for operations traversing multiple regions of the national airspace system.

From an approach standpoint, the inhomogeneous Markov technique developed in Chapter 2 was successful in illustrating the regional uncertainty characteristics of the RUC20 product. Realizations of the model clearly illustrate the impact of the jet-stream on a data assimilation model relying heavily on aircraft reports, with the RUC 20 exhibiting a slight positive bias along routes where aircraft leverage the strong tailwind provided by the jet stream. The converse case is also evident for the first time, with the RUC20 displaying a slight negative bias along routes where aircraft actively avoid strong jet stream headwinds. This result demonstrates a need for models that capture regional forecast uncertainty characteristics, as the uncertainty characteristics of the RUC forecast product as an example are inexorably driven by the air route. Regarding the issue of error scale, the inhomogeneous Markov approach was unable to address this issue directly; on its face, the reasoning should be rather plain to the reader as well. Any model hoping to capture an error effect within the RUC20 will require co-temporal measurements of equal or lesser than scale compared to the effect itself. In other words, the only real solution to this issue is to increase the sensor network density. Although the inhomogeneous Markov approach was unable to solve an issue inherent to the problem itself, it is completely adaptable to increased sensor network density. As the sensor network (i.e.: density of aircraft wind reports) increases, the longitudinal scale of the model can be increased to capture the effects of smaller scale forecast features.

From an application standpoint, an issue with the forecast uncertainty model was briefly addressed in Section 3.4.3. Initializing the uncertainty model with a variety of initial distributions provides very little impact on the realizations of the model itself. Upon closer inspection, this phenomenon is explained by data density issues. Of the 2,797 flight paths between KSEA and KATL comprising the data set

for this airport origin-destination pair, only 217 flights directly contributed to the uncertainty model constructed along sample route 1. Of these 217 flights, very few exhibited consistently large forecast error measurements. Accordingly, the transition probabilities tend to drive most fringe measurements towards the center when the model is actually realized. Though detrimental towards illustrating the possible uses of the stochastic algorithm as the solution space becomes artificially narrow, the approach itself is still valid and the models themselves can be improved with access to a better data set. Currently, U.S. air carriers only contribute a fraction of collected meteorological data reports to the ACARS program. The author’s methodology could very easily be utilized to construct high resolution model along popular routes within a carrier’s route structure using the wealth of privately held data.

Current FMS RTA mode implementations are purely deterministic. Research efforts surrounding fuel-optimal RTA functionality have approached the issue from a large-scale, linearized model approach for ground-based solutions, and development of on-board solutions have approached RTA functionality given forecast uncertainty from a purely adherence-driven standpoint utilizing a “worst case scenario” approach. Studies have repeatedly shown that an on-board system is preferred for a variety of reasons, however, any approach combining fuel optimization strategies utilizing large scale mixed integer programming techniques is immediately at odds with the technology available in the cockpit. A clear need for an algorithm capable of including wind speed forecast uncertainty information in fuel-centric RTA speed while being computationally tractable with an on-board system is plainly evident.

The stochastic algorithm developed in Chapter 3 succeeds in bridging the gap between fuel optimization approaches developed in the literature and the constraints imposed by a lack of compute power. In most cases, the quadratic search technique allows the search heuristic to converge within one iteration (though an additional iteration is required to verify the solution). Moreover, the algorithm demonstrates

fuel savings on the order of 50-100 *lbs* of fuel for transcontinental RTA operations when compared to standard dead-band control approaches. Computationally, the algorithm compares very favorably performance-wise to the large scale optimization approaches described in the literature. Generally, mixed integer programs solve in polynomial time. The algorithm proposed by the author solves in better than linear time, consistently providing solutions in fractional power time meaning performance improves as the solution space increases in size. However, though a very interesting result from a computational complexity standpoint, it is important to note that performance gains exhibited by algorithms with better than linear time to solutions are only really evident once the solution space reaches a certain size. For the solution space of 10,000 scenarios proposed in this effort, the algorithm performs comparably to a linear time approach.

4.1 Contributions to the Field

This research effort provides two significant contributions to the field: a forecast uncertainty modeling methodology capable of capturing regional forecast uncertainty characteristics, and a stochastic fuel-optimal RTA algorithm which is computationally tractable given the limited resources of on-board systems. The uncertainty model developed in this effort clearly demonstrates the need for regional or route specific models by plainly illustrating the inherent bias of the RUC20 model due to jet stream effects, and the RTA algorithm provides a compact approach to fuel-optimal RTA performance with an algorithm that performs better than its large-scale counterparts from a computational complexity standpoint given a much smaller implementation footprint.

4.2 *Future Work*

Both the wind forecast uncertainty model and the required time of arrival algorithm warrant further exploration. To begin, the wind forecast uncertainty modeling methodology relied on the fact that aircraft flight track data was concentrated along specific routes through the airspace, excluding flights from the considered data set which did not closely follow the proposed sample route. In order to generate a Markov model utilizing fixed transition probability matrices, this characteristic of the model is an unfortunate necessity. Expanding the analysis methodology to instead consider a variable route and thus variable transition probabilities would allow for a more flexible approach. As a proposed solution, Kriging could be used to estimate the transition probabilities as a function of unique route and neighboring data points rather than tying the transition probabilities to a predefined route.

It may have become evident to the reader that the RTA algorithm produces cruise speed recommendations that often lie in close proximity to the nominal Mach number required for RTA adherence. If one further considers that modern flight management systems can only be expected to accept Mach inputs with precision to five thousandths of a Mach number, the entire algorithm could be reformulated to consider Mach number as a discrete rather than continuous variable, with Mach number discretized in 0.005 Mach increments. Doing so would severely restrict the solution space allowing for the consideration of additional variables such as the location of the decision waypoint. Rather than modeling the decision waypoint as a constraint, the acceptable width of the arrival time distribution could be constrained based on carrier RTA adherence risk tolerance, and the location of the decision waypoint explored to find the combination of stage one Mach number and decision waypoint location which minimizes fuel burn while ensuring that arrival time uncertainty remains bounded.

APPENDIX A

AIRCRAFT ROUTES

A.1 Route 2: KATL to KSEA

Table 9: Route used to construct uncertainty model for flights between KATL and KSEA.

Waypoint	Latitude ($^{\circ}N$)	Longitude ($^{\circ}E$)	Length (nm)	Heading ($^{\circ}$)
BNA	36.1369608	-86.6847717	243.926	-46.9
STL	38.8606889	-90.4823681	278.852	-58.6
OVR	41.1672222	-95.7369444	363.702	-60.0
RAP	43.9760292	-103.0123419	263.352	-63.4
BIL	45.8085592	-108.6246475	146.739	-69.7
HLN	46.6068208	-111.9534753	159.893	-70.0
MLP	47.4569031	-115.646045	80.827	-84.7
GEG	47.5649444	-117.6268889	N/A	N/A

KATL to KSEA

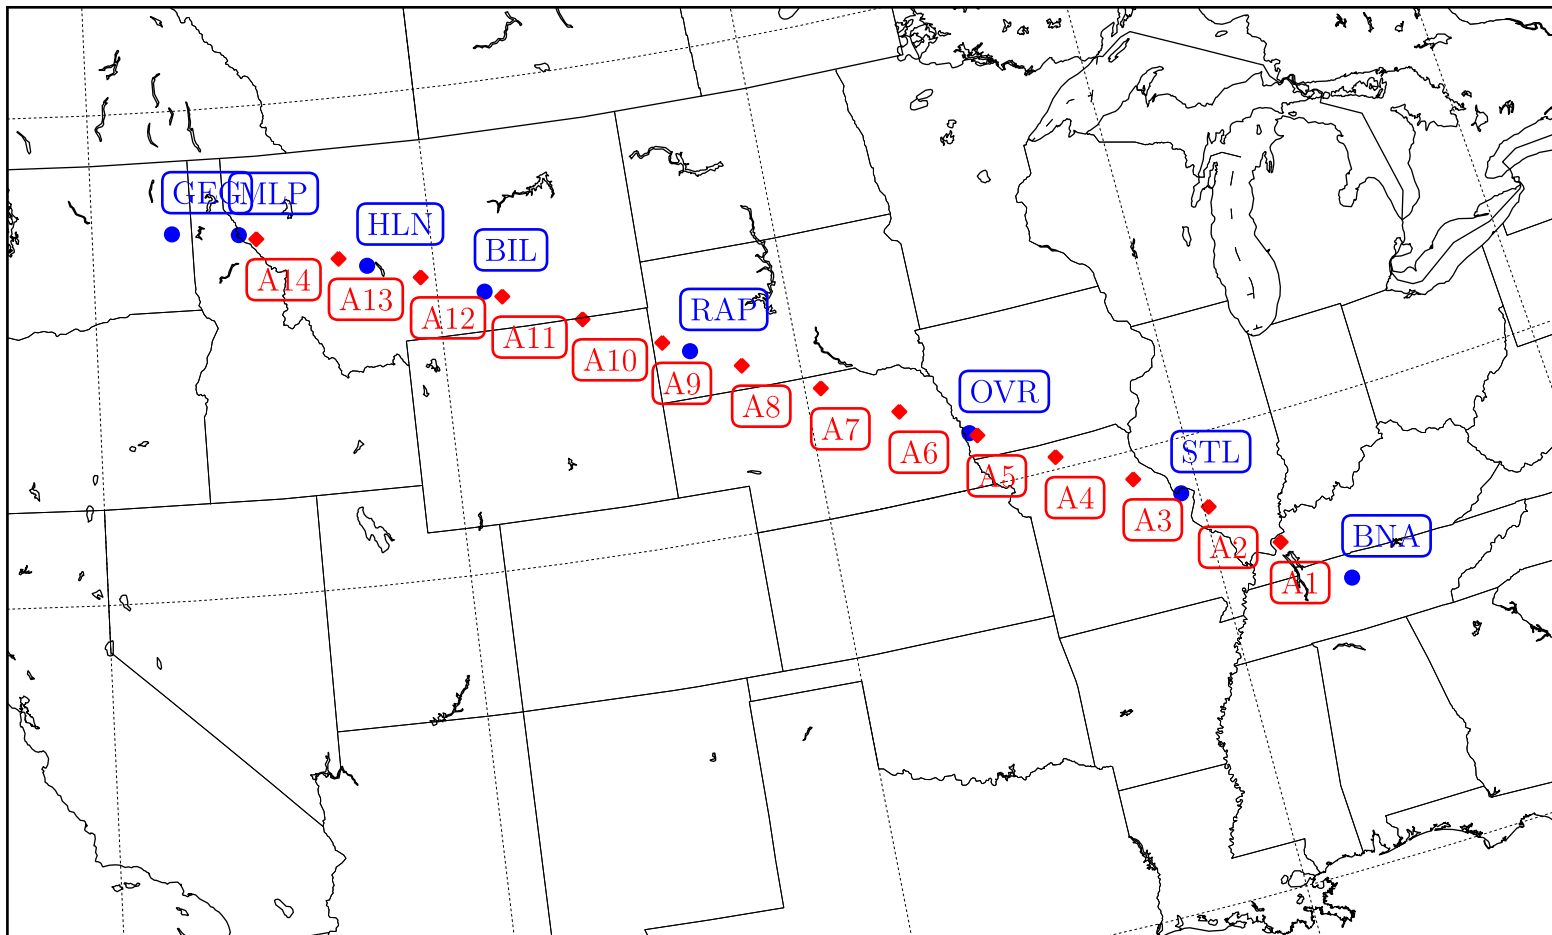


Figure 34: KATL to KSEA route with standard and artificial waypoints.

A.2 Route 3: KLAX to KATL

Table 10: Route used to construct uncertainty model for flights between KLAX and KATL.

Waypoint	Latitude ($^{\circ}N$)	Longitude ($^{\circ}E$)	Length (nm)	Heading ($^{\circ}$)
TRM	33.6280833	-116.1601944	353.070	80.3
SJN	34.4240367	-109.14352	206.966	77.4
ACH	35.1117061	-105.0399289	164.482	86.5
PNH	35.235	-101.6991667	101.585	85.7
SYO	35.3452778	-99.6352778	99.442	89.0
IRW	35.3586111	-97.6091667	163.776	88.4
FSM	35.3883333	-94.2713889	212.016	94.8
MEM	35.0151111	-89.9832222	N/A	N/A

KLAX to KATL

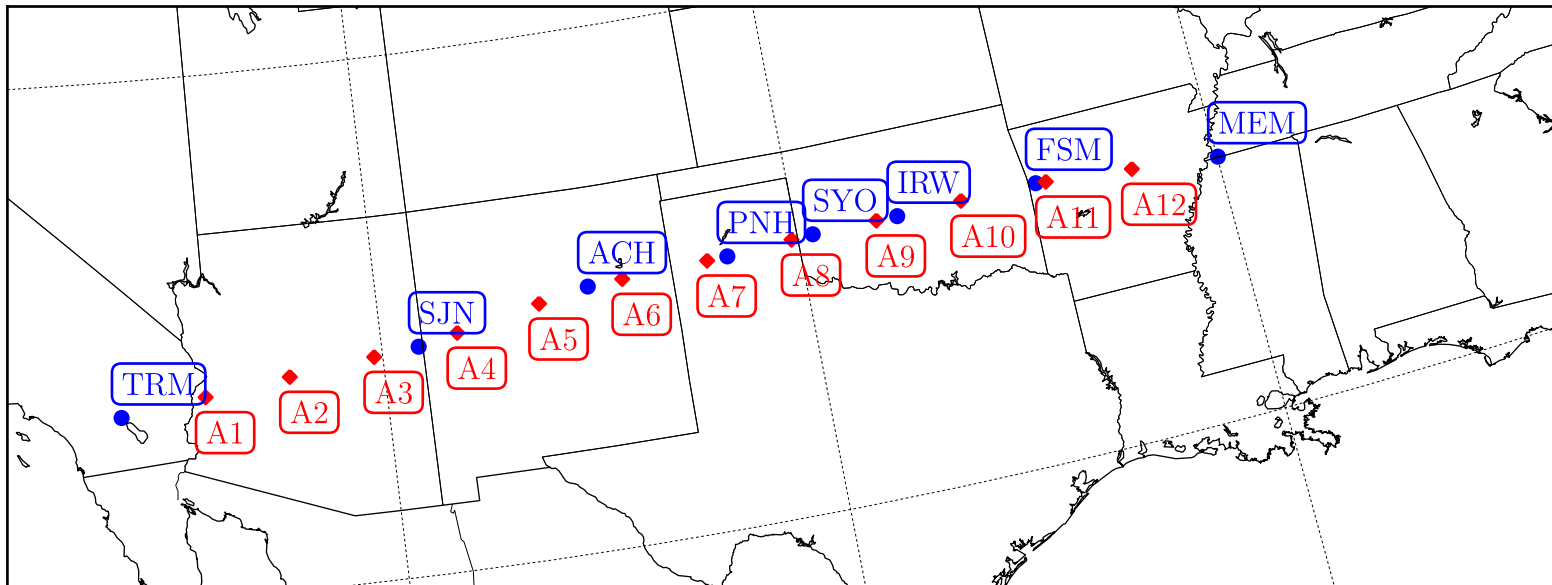


Figure 35: KLAX to KATL route with standard and artificial waypoints.

A.3 Route 4: KATL to KLAX

Table 11: Route used to construct uncertainty model for flights between KATL and KLAX.

Waypoint	Latitude ($^{\circ}N$)	Longitude ($^{\circ}E$)	Length (nm)	Heading ($^{\circ}$)
VUZ	33.6701389	-86.8998333	269.684	-75.6
LIT	34.6776728	-92.1805283	270.586	-79.8
IRW	35.3585892	-97.6092336	200.984	-90.9
PNH	35.2350628	-101.6990328	164.492	-91.6
ACH	35.1116667	-105.04	87.568	-92.2
ABQ	35.0437956	-106.8163119	115.339	-91.7
ZUN	34.9657533	-109.1545094	165.013	-94.5
DRK	34.7025564	-112.4803492	167.084	-101.3
TNP	34.1122222	-115.7699167	N/A	N/A

KATL to KLAX

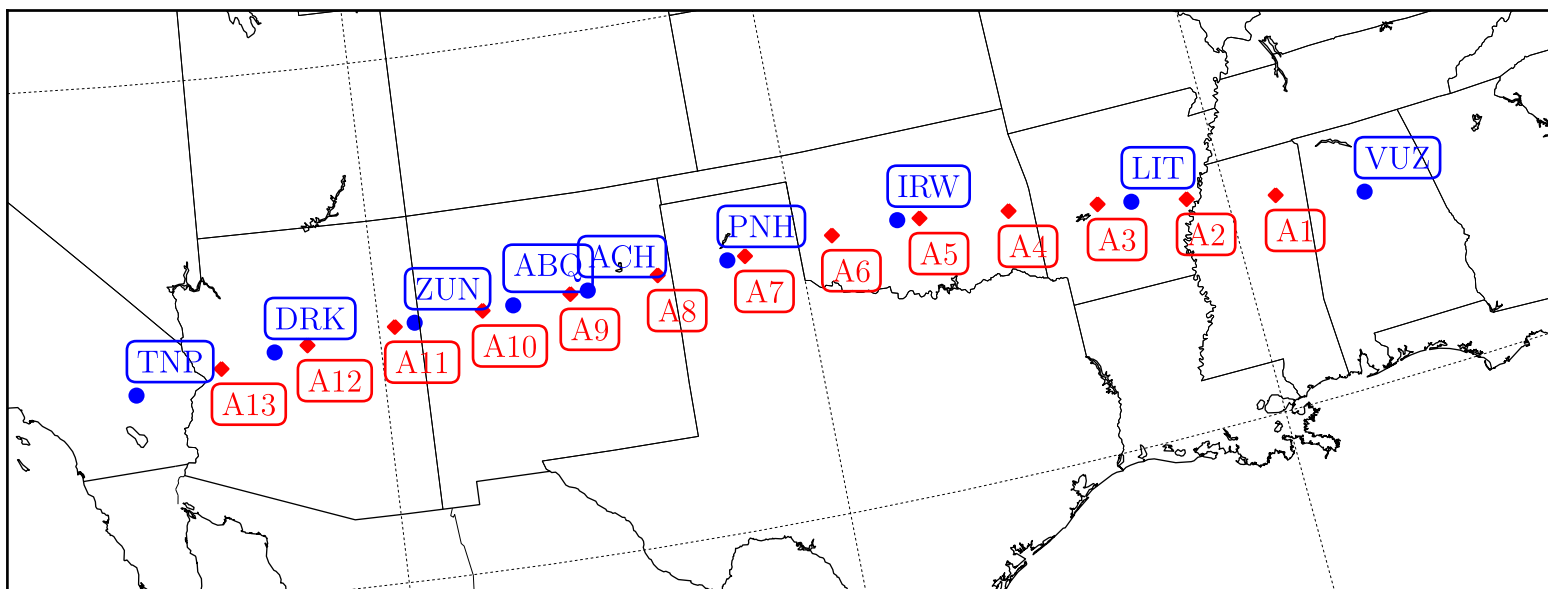


Figure 36: KATL to KLAX route with standard and artificial waypoints.

A.4 Route 5: KEWR to KATL

Table 12: Route used to construct uncertainty model for flights between KEWR and KATL.

Waypoint	Latitude ($^{\circ}N$)	Longitude ($^{\circ}E$)	Length (nm)	Heading ($^{\circ}$)
PTW	40.2222364	-75.5602589	78.708	-123.2
EMI	39.4950075	-76.9785719	65.857	-140.7
CSN	38.6412022	-77.8654994	73.561	-126.7
MOL	37.9005247	-79.1068892	279.546	-132.1
ODF	34.6958611	-83.2976667	N/A	N/A

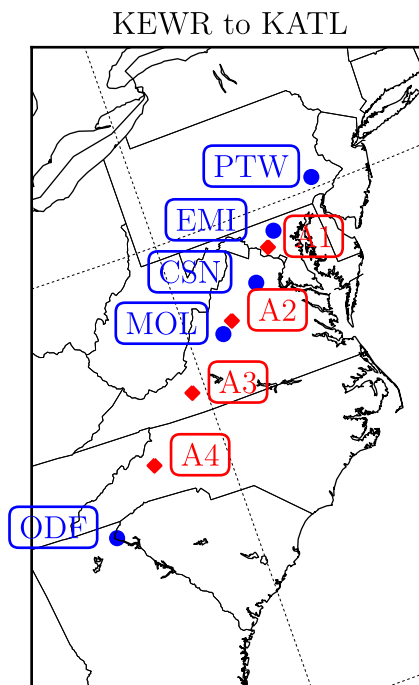


Figure 37: KEWR to KATL route with standard and artificial waypoints.

A.5 Route 6: KATL to KEWR

Table 13: Route used to construct uncertainty model for flights between KATL and KEWR.

Waypoint	Latitude ($^{\circ}N$)	Longitude ($^{\circ}E$)	Length (nm)	Heading ($^{\circ}$)
SPA	35.0336389	-81.927	113.129	57.0
GSO	36.0456919	-79.9763753	136.424	48.7
FAK	37.5285	-77.8282222	87.241	35.7
OTT	38.7058611	-76.74475	78.824	42.0
DQO	39.6781389	-75.6070833	47.227	42.9
ARD	40.2533333	-74.9076111	N/A	N/A

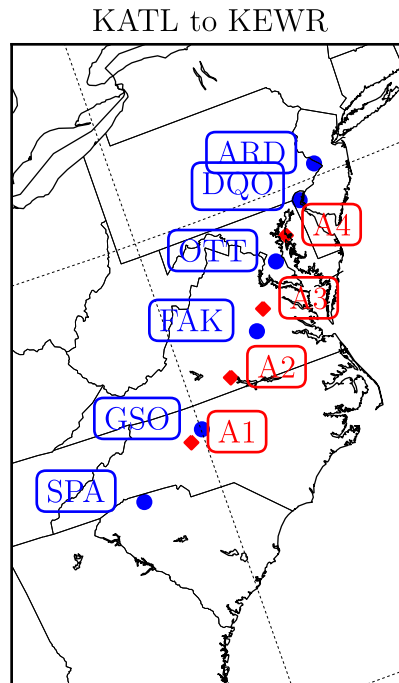


Figure 38: KATL to KEWR route with standard and artificial waypoints.

REFERENCES

- [1] BEALE, E. M. L., “On Minimizing a Convex Function Subject to Linear Inequalities,” *Journal of the Royal Statistical Society, Series B*, vol. 17, no. 2, pp. 173–184, 1955.
- [2] BENDERS, J. F., “Partitioning Procedures for Solving Mixed-Variables Programming Problems,” *Numerische Mathematik*, vol. 4, pp. 238–252, 1962.
- [3] BENJAMIN, S. G., “An Isentropic Mesoalpha-Scale Analysis System and Its Sensitivity to Aircraft and Surface Observations,” *Monthly Weather Review*, vol. 117, pp. 1586–1603, July 1989.
- [4] BENJAMIN, S. G., BREWSTER, K. A., BRÜMMER, R., JEWETT, B. F., SCHLATTER, T. W., SMITH, T. L., and STAMUS, P. A., “An Isentropic Three-Hour Data Assimilation System Using ACARS Aircraft Observations,” *Monthly Weather Review*, vol. 119, pp. 888–906, Apr. 1990.
- [5] BENJAMIN, S. G., BROWN, J. M., BRUNDAGE, K. J., DÉVÉNYI, D., GEORG, KIM, D., SCHWARTZ, B. E., SMIRNOVA, T. G., SMITH, T. L., WEYGANDT, S. S., and MANIKIN, G. S., “NWS Technical Procedures Bulletin No. 490: RUC20 - The 20-km Version of the Rapid Update Cycle,” tech. rep., NOAA/OAR Forecast Systems Laboratory, Boulder, CO, May 2002.
- [6] BENJAMIN, S. G., BRUNDAGE, K. J., MILLER, P. A., SMITH, T., GRELL, G. A., KIM, D., BROWN, J. M., and SCHLATTER, T. W., “The Rapid Update Cycle at NMC. Preprints,” in *10th Conference on Numerical Weather Prediction, Portland, Oregon, 18-22 Jul. 1994*.
- [7] BENJAMIN, S. G., DÉVÉNYI, D., WEYGANDT, S. S., BRUNDAGE, K. J., BROWN, J. M., GRELL, G. A., KIM, D., SCHWARTZ, B. E., SMIRNOVA, T. G., SMITH, T. L., and MANIKIN, G. S., “An Hourly Assimilation–Forecast Cycle: The RUC,” *Monthly Weather Review*, vol. 132, pp. 495–518, Feb. 2004.
- [8] BENJAMIN, S. G., GRELL, G. A., BROWN, J. M., SMIRNOVA, T. G., and BLECK, R., “Mesoscale Weather Prediction with the RUC Hybrid Isentropic–Terrain-Following Coordinate Model,” *Monthly Weather Review*, vol. 132, pp. 473–494, Feb. 2004.
- [9] BENJAMIN, S. G., SCHWARTZ, B. E., and COLE, R. E., “Accuracy of ACARS Wind and Temperature Observations Determined by Collocation,” *Weather and Forecasting*, vol. 14, pp. 1032–1038, Dec. 1999.

- [10] BIRGE, J. R. and LOUVEAUX, F. V., “A Multicut Algorithm for Two-Stage Stochastic Linear Programs,” *European Journal of Operational Research*, vol. 34, no. 3, pp. 384–392, 1988.
- [11] BISIAUX, M., COX, M. E., FORRESTER, D. A., and STOREY, J. T., “Possible Improvements in Meteorology for Aircraft Navigation,” *Journal of Navigation*, vol. 36, pp. 54–63, Jan. 1983.
- [12] BLIN, K., AKIAN, M., BONNANS, F., HOFFMAN, E., MARTINI, C., and ZAGHAL, K., “A Stochastic Conflict Detection Model Revisited,” in *AIAA Guidance, Navigation and Control AIAA Guidance, Navigation, and Control Conference and Exhibit*, Reston, VA, 14-17 Aug. 2000.
- [13] BURROWS, J. W. and CHAKRAVARTY, A., “Time-Controlled Aircraft Guidance in Uncertain Winds and Temperatures,” in *American Control Conference*, Seattle, WA, 608 Jun., 1984.
- [14] CHALOULOS, G. and LYGEROS, J., “Effect of wind correlation on aircraft conflict probability,” *Journal of Guidance, Control, and Dynamics*, vol. 30, pp. 1742–1752, Nov. 2007.
- [15] COLE, R., RICHARD, C., KIM, S., and BAILY, ., “An Assessment of the 60 km Rapid Update Cycle (RUC) with Near Real-Time Aircraft Reports,” tech. rep., Project Report NASA/A-1, MIT Lincoln Laboratory, Lexington, MA, 15 Jul. 1998.
- [16] COLE, R. E., GREEN, S., JARDIN, M., SCHWARTZ, B. E., and BENJAMIN, S. G., “Wind Prediction Accuracy for Air Traffic Measurement Decision Support Tools,” in *3 rd USA/Europe Air Traffic Management R&D Seminar, Napoli, Italy*, 13-16 Jun. 2000.
- [17] COLE, R. E., GREEN, S. M., and JARDIN, M. R., “Improving RUC-1 Wind Estimates by Incorporating Near-Real-Time Aircraft Reports,” *Weather and Forecasting*, vol. 15, pp. 447–460, Aug. 2000.
- [18] COLE, R. E. and KIM, S., “A study of Time-to-Fly Estimates for RUC and ITWS Winds,” in *Ninth Conference on Aviation Range, and Aerospace Meteorology (ARAM)*, Orlando, FL, 2000.
- [19] DANTZIG, G. B. and WOLFE, P., “Decomposition Principle for Linear Programs,” *Operations Research*, vol. 8, pp. 101–111, 1960.
- [20] DANTZIG, G. B. and WOLFE, P., “The Decomposition Algorithm for Linear Programming,” *Econometrica*, vol. 29, no. 4, pp. 767–778, 1961.
- [21] DANTZIG, G. B., “Linear Programming Under Uncertainty,” *Management Science*, vol. 1, pp. 197–206, Apr. - Jul. 1955.

- [22] DE MENORVAL, J.-L., LUCAS, F., and GERASIMOVIC, D., “Method and Device to Assist in the Guidance of an Airplane.” US2008/0228333 A1, Mar. 11 2008.
- [23] DEJONGE, M. K., “Required Time of Arrival (RTA) Control System.” United States Patent 5,121,325, Jun. 9 1993.
- [24] DURAND, N. and ALLIOT, J.-M., “Optimal Resolution of En Route Conflicts,” in *1st USA/Europe Air Traffic Management R&D Seminar, Toulouse, France*, 17-20 Jun. 1997.
- [25] EARTH SYSTEMS RESEARCH LABORATORY, NOAA, “RUC FAQ.” <http://ruc.noaa.gov/RUC.faq.html>.
- [26] EUROCONTROL, “Base of Aircraft Data (BADA).” <http://www.eurocontrol.int/services/bada>.
- [27] FAA, “Concept of Operations for the Next Generation Air Transportation System,” Tech. Rep. v. 2.0, Joint Planning and Development Office, 2008.
- [28] FAA, “FAA Aerospace Forecast: Fiscal Years 2008-2025,” tech. rep., Federal Aviation Administration, 2008.
- [29] FAA, “Next Generation Air Transportation System Integrated Work Plan: A Functional Outline,” Tech. Rep. v. 1.0, Joint Planning and Development Office, 2008.
- [30] FEDERAL AVIATION ADMINISTRATION, “Air Traffic Bulletin: 2003-5.” http://www.faa.gov/air_traffic/publications/atpubs/atbarc/03-5.htm.
- [31] FEDERAL AVIATION ADMINISTRATION, “Federal Aviation Regulations.” Part 91, §91.179.
- [32] GOSNER, J. M. and KOMINEK, R. J., “Apparatus and Method for Controlling an Aircraft Performance Calculator to Achieve Time-Constrained Navigation.” United States Patent 5,408,413, Apr. 18 1995.
- [33] GREEN, S. M. and VIVONA, R. A., “Field Evaluation of Descent Advisor Trajectory Prediction Accuracy,” in *AIAA Guidance, Navigation and Control AIAA Guidance, Navigation, and Control Conference, San Diego, CA*, 29-31 Jul. 1996.
- [34] HAMILL, T. M. and WHITAKER, J. S., “Model Errors in Ensemble Forecasts: The Structure of Error from Unrepresented Scales,” *Monthly Weather Review*, vol. 133, no. 11, pp. 3132–3147, 2005.
- [35] HARALDSDOTTIR, A., SCHARL, J., KING, J., SCHOEMIG, E. G., and BERGE, M. E., “Analysis of Arrival Management Performance with Aircraft Required Time of Arrival Capabilities,” in *26th International Congress of the Aeronautical Sciences, Anchorage, Alaska*, 14-19 Sep. 2008.

- [36] HOFFMAN, R. N. and KALNAY, E., “Lagged Averaged Forecasting, an Alternative to Monte Carlo Forecasting,” *Tellus A: Dynamic Meteorology and Oceanography*, vol. 35A, pp. 100–118, Mar. 1983.
- [37] HOLLINGSWORTH, A. and LÖNNBERG, P., “The Statistical Structure of Short-Range Forecast Error as Determined from Radiosone Data. Part I: The Wind Field,” *Tellus A: Dynamic Meteorology and Oceanography*, vol. 38A, no. 2, pp. 111–136, 1986.
- [38] HOWARD, R. A., *Dynamic Probabilistic Systems, Volume I: Markov Models*. New York, New York: John Wiley and Sons, 1971.
- [39] HU, J., PRANDINI, M., and SASTRY, S., “Aircraft Conflict Detection in the Presence of Spatially Correlated Wind Field,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 6, pp. 326–340, Sept. 2005.
- [40] IATA, “Fuel Action Plan: Guidance Material and Best Practices for Fuel and Environmental Management,” tech. rep., International Air Transport Association, 2004.
- [41] ICAO, “Manual of the ICAO Standard Atmosphere, extended to 80 kilometers (262 500 feet). Third Edition,” Tech. Rep. Doc 7488/3, International Civil Aviation Organization, Montreal, Quebec, Canada, 1993.
- [42] IOSIFESCU, M., *Finite Markov Processes and their Applications*. John Wiley and Sons, 1980.
- [43] KARMARKAR, N., “A New Polynomial-Time Algorithm for Linear Programming,” in *Proceedings of the 16th Annual ACM Symposium on Theory of Computing (1984)* (ACM, ed.), (New York), pp. 302–311, 1984.
- [44] KELLEY, J. E., “The Cutting Plane Method for Solving Convex Programs,” *Journal of the SIAM*, vol. 8, pp. 703–712, 1960.
- [45] KEYWEGT, A. D., SHAPIRO, A., and DE MELLO, T. H., “The sample average approximation method for stochastic discrete optimization,” *SIAM Journal on Optimization*, vol. 12, no. 2, pp. 479–502, 2001.
- [46] KHACHIYAN, L. G., “Polynomial Algorithms in Linear Programming,” *USSR Computational Mathematics and Mathematical Physics*, vol. 20, no. 1, pp. 53–72, 1980.
- [47] KLEE, V. and MINTY, G. J., “How Good is the Simplex Algorithm?,” Tech. Rep. TR-22, Washington University, Seattle, Department of Mathematics, 1970.
- [48] KRIGE, D. G., “A Statistical Approach to Some Basic Mine Valuation Problems on the Witwatersrand,” *Journal of the Chemical, Metallurgical and Mining Society of South Africa*, vol. 52, no. 6, pp. 119–139, 1951.

- [49] KUCHAR, J. K. and YANG, L. C., “A review of Conflict Detection and Resolution Modelling Methods,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 1, pp. 179–189, Dec. 2000.
- [50] LANCIANI, A. and SALVATI, M., “Spatial Interpolation of Surface Weather Observations in Alpine Meteorological Services,” Tech. Rep. FORLAPS Technical Report, 2, Università degli Studi di Trento, Dipartimento di Ingegneria Civile e Ambientale, Trento, Italy, Mar. 2008.
- [51] LEE, A. G., WEYGANDT, S. S., SCHWARTZ, B., and MURPHEY, J. R., “Performance of Trajectory Models with Wind Uncertainty,” in *AIAA Modeling and Simulation Technologies Conference, Chicago, IL*, 10-13 Aug. 2009.
- [52] LEIDNER, S. M., STAUFFER, D. R., and SEAMAN, N. L., “Improving Short-Term Numerical Weather Prediction in the California Coastal Zone by Dynamic Initialization of the Marine Boundary Layer,” *Monthly Weather Review*, vol. 129, pp. 275–294, Feb. 2001.
- [53] LEWIS, J. M., “Ooishi’s Observation: Viewed in the Context of Jet Stream Discovery,” *Bulletin of the American Meteorological Society*, vol. 84, pp. 357–369, Mar. 2003.
- [54] LINDROTH, J., SHAPIRO, A., and WRIGHT, S., “The Empirical Behavior of sampling Methods for Stochastic Programming,” *Annals of Operations Research*, vol. 142, pp. 215–241, Feb. 2006.
- [55] LORENC, A. C., BELL, R. S., and MACPHERSON, B., “The Meteorological Office Analysis Correction Sata Assimilation Scheme,” *Quarterly Journal of the Royal Meteorological Society*, vol. 117, pp. 58–89, Jan. 1991.
- [56] LOWTHER, M. B., “En Route Speed Optimization for Continuous Descent Arrival,” Master’s thesis, Georgia Institute of Technology, May 2008.
- [57] LU, C., YUAN, H., SCHWARTZ, B. E., and BENJAMIN, S. G., “Short-Range Numerical Weather Prediction Using Time-Lagged Ensembles,” *Weather and Forecasting*, vol. 22, no. 3, pp. 580–595, 2007.
- [58] LYGEROS, J. and PRANDINI, M., “Aircraft and Weather Models for Probabilistic Collision Avoidance in Air Traffic Control,” in *41st IEEE International Conference on Decision and Control, Las Vegas, Nevada*, 10 - 13 Dec. 2002.
- [59] LYMPEROPOULOS, I. and LYGEROS, J., “Sequential Monte Carlo Methods for Multi-Aircraft Trajectory Prediction in Air Traffic Management,” *International Journal of Adaptive Control and Signal Processing*, vol. 24, pp. 830–849, Oct. 2010.
- [60] MACPHERSON, B., “Dynamic Initialization by repeated Insertion of Data,” *Quarterly Journal of the Royal Meteorological Society*, vol. 117, pp. 965–991, Jul. 1991.

- [61] MONDOLONI, S., “Aircraft Trajectory Prediction Errors: Including a Summary of Error Sources and Data,” tech. rep., Federal Aviation Administration, July 2006.
- [62] MONDOLONI, S., “A Multiple-Scale Model of Wind-Prediction Uncertainty and Application to Trajectory Prediction,” in *6th AIAA Aviation Technology, Integration and Operations Conference (ATIO)*, Wichita, Kansas, 25 - 27 Sept. 2006.
- [63] MONDOLONI, S., PAGLIONE, M., and GREEN, S., “Trajectory Modeling Accuracy for Air Traffic Management Decision Support Tools,” in *ICAS 2002 Congress, Toronto, Canada*, September 2002.
- [64] NATIONAL CENTERS FOR ENVIRONMENT PROTECTION, “NCEP Grid 252.” <http://www.nco.ncep.noaa.gov/pmb/docs/on388/grids/grid252.gif>.
- [65] NATIONAL WEATHER SERVICE, “Jet Stream - Online School for Weather.” <http://www.srh.noaa.gov/jetstream/global/jet.htm>.
- [66] NICHOLLS, S., “Aircraft Observations of the Ekman Layer during the Joint Air-Sea Interaction Experiment,” *Quarterly Journal of the Royal Meteorological Society*, vol. 111, pp. 391–346, Apr. 1985.
- [67] NORRIS, J. R., *Markov Chains*. Cambridge, United Kingdom: Cambridge University Press, 1997.
- [68] OSTWALD, P., “Impacts of ATC Related Maneuvers on Meeting a Required Time of Arrival,” in *25th Digital Avionic Systems Conference, Portland, OR*, 15-18 Oct. 2006.
- [69] PAIELLI, R. A. and ERZBERGER, H., “Conflict Probability Estimation for Free Flight,” tech. rep., NASA Technical Memorandum 110411, Oct. 1996.
- [70] PRINCE, G. V. and MACPHERSON, A. K., “A numerical weather forecasting method using cubic splines on a variable mesh,” *Journal of Applied Meteorology*, vol. 12, pp. 1102–1113, Oct. 1973.
- [71] RIEHL, H., “Jet Stream in Upper Troposphere and Cyclone Formation,” *Transactions, American Geophysical Union*, vol. 29, no. 2, pp. 175–186, 1948.
- [72] ROBERSON, B., “Fuel Conservation Strategies: Cost Index Explained,” *Aero Magazine*, vol. Quarter 2, 2007.
- [73] ROSSBY, C. G., “On the distribution of angular velocity in gaseous envelopes under the influence of On the distribution of angular velocity in gaseous envelopes under the influence of On the Distribution of Angular Velocity in Gaseous Envelopes Under the Influence of Large-Scale Horizontal Mixing Processes,” *Bulletin of the American Meteorological Society*, vol. 28, pp. 53–68, 1947.

- [74] RUMBO, J. R., JACKSON, M. R., and O’LAUGHLIN, B. E., “Aircraft Control System for Reaching a Waypoint at a Required Time of Arrival.” United States Patent 6,507,782, Jan. 14 2003.
- [75] RUSZCZYNSKI, A., “A Regularized Decomposition Method for Minimizing a Sum of Polyhedral Functions,” *Mathematical Programming*, vol. 35, no. 3, pp. 309–333, 1986.
- [76] SCHARL, J., HARALDSDOTTIR, A., J. KING, R. S., and WICHMAN, K. D., “A Fast-Time Required Time of Arrival Model for Analysis of 4D Arrival Management Concepts,” in *AIAA Modeling and Simulation Technologies Conference, Honolulu, Hawaii*, 8-21 Aug. 2008.
- [77] SCHWARTZ, B. and BENJAMIN, S. G., “A Comparison of temperature and Wind Measurements from ACARS-Equipped Aircraft and Rawinsondes,” *Weather and Forecasting*, vol. 10, pp. 528–544, Sept. 1995.
- [78] SCHWARTZ, B. E., BENJAMIN, S. G., GREEN, S. M., and JARDIN, M. R., “Accuracy of RUC-1 and RUC-2 Wind and Aircraft Trajectory Forecasts by Comparison with ACARS Observations,” *Weather and Forecasting*, vol. 15, pp. 313–326, June 2000.
- [79] SHAPIRO, A. and DE MELLO, T. H., “On the Rate of Convergence of Optimal Solutions of Monte Carlo Approximations of Stochastic Programs,” *SIAM Journal on Optimization*, vol. 11, no. 1, pp. 70–86, 2000.
- [80] SHAPIRO, A., DENTCHEVA, D., and RUSZCZYŃSKI, A., *Lectures on Stochastic Programming: Modeling and Theory*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2009.
- [81] SHEPARD, D., “A Two-Dimensional Interpolation Function for Irregularly-Spaced Data,” in *Proceedings of the 1968 23rd ACM National Conference*, (New York, New York), pp. 517–524, Association for Computing Machinery, 1968.
- [82] SMEDT, D. D. and BERZ, G., “Study of the Required Time of Arrival Function of Current FMS in an ATM Context,” in *6th Digital Avionics Systems Conference, Dallas, Texas*, 21-25 Oct. 2007.
- [83] STAFF MEMBERS, “On the Circulation of the Atmosphere in Middle Latitudes,” *Bulletin of the American Meteorological Society*, vol. 28, pp. 255–280, 1947.
- [84] STAUFFER, D. R. and SEAMAN, N. L., “Multiscale Four-Dimensional Data Assimilation,” *Journal of Applied Meteorology*, vol. 33, pp. 416–434, Mar. 1994.
- [85] TINO, C. P. and CLARKE, J.-P. B., “Systems And Methods Providing A Fuel-Efficient RTA Implementation With Uncertain Winds.” United States Patent Application 13/624,771, 21 Sept. 2012.

- [86] TINO, C. P., REN, L., and CLARKE, J.-P. B., “Wind Forecast Error and Trajectory prediction for En-Route Scheduling,” in *AIAA Guidance, Navigation and Control AIAA Guidance, Navigation, and Control Conference and Exhibit, Chicago Illinois*, 10 - 13 August 2009.
- [87] VINCENTY, T., “Direct and inverse Solutions of Geodesics on the Ellipsoid with Application of Nested Equations,” *Survey Review*, vol. 22, pp. 88–93, April 1975.
- [88] VISINTINI, A. L., GLOVER, W., LYGEROES, J., and MACIEJOWSKI, J., “Monte Carlo Optimization for Conflict Resolution in Air Traffic Control,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, pp. 470–482, Dec. 2006.
- [89] WICHMAN, K. D., CARLSSON, G., and LINDBERG, L. G., “Flight Trials: ”Runway-to-Runway” Required Time of Arrival Evaluations for Time-Based ATM Environment - Final Results,” in *AIAA Guidance, Navigation and Control AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, California*, 5–8 Aug. 2002.
- [90] WILLIAMS, D. H. and GREEN, S. M., “Airborn Four-Dimensional Flight Management in a Time-Based Air Traffic Control Environment,” tech. rep., NASA Technical Memorandum 4249, 1991.
- [91] ZHENG, Q. M. and ZHAO, Y., “Modeling Wind Uncertainties for Stochastic Trajectory Synthesis,” in *11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, Virginia Beach, VA*, 20 - 22 Sept. 2011.