

ESTIMATION OF POTENTIAL AIRCRAFT FUEL BURN REDUCTION IN CRUISE VIA SPEED AND ALTITUDE OPTIMIZATION STRATEGIES

Jonathan A. Lovegren and R. John Hansman

This report is based on the Master of Science Thesis of Jonathan A. Lovegren submitted to the Department of Aeronautics and Astronautics in partial fulfillment of the requirements for the degree of Master of Science in Aeronautics and Astronautics at the Massachusetts Institute of Technology.

Report No. ICAT-2011-03 February 2011

MIT International Center for Air Transportation (ICAT) Department of Aeronautics & Astronautics Massachusetts Institute of Technology Cambridge, MA 02139 USA

Quantification of Fuel Burn Reduction in Cruise Via Speed and Altitude Optimization Strategies

by

Jonathan A. Lovegren and R. John Hansman

Abstract

Environmental performance has become a dominant theme in all transportation sectors. As scientific evidence for global climate change mounts, social and political pressure to reduce fuel burn and CO_2 emissions has increased accordingly, especially in the rapidly growing aviation industry. Operational improvements offer the ability to increase the performance of any aircraft immediately, by simply changing how the aircraft is flown. Cruise phase represents the largest portion of flight, and correspondingly the largest opportunity for fuel burn reduction.

This research focuses on the potential efficiency benefits that can be achieved by improving the cruise speed and altitude profiles operated by flights today. Speed and altitude are closely linked with aircraft performance, so optimizing these profiles offers significant fuel burn savings. Unlike lateral route optimization, which simply attempts to minimize the distance flown, speed and altitude changes promise to increase the efficiency of aircraft throughout the entire flight.

Flight data was collected for 257 flights during one day of domestic US operations. A process was developed to calculate the cruise fuel burn of each selected flight, based on aircraft performance data obtained from Piano-X and atmospheric data from NOAA. Improved speed and altitude profiles were then generated for each flight, representing various levels of optimization. Optimal cruise climbs and step climbs of 1,000 and 2,000 ft were analyzed, along with optimal and LRC speed profiles.

Results showed that a maximum fuel burn reduction of 3.5% is possible in cruise given complete altitude and speed optimization; this represents 2.6% fuel reduction system-wide, corresponding to 300 billion gallons of jet fuel and 3.2 million tons of CO_2 saved annually. Flights showed a larger potential to improve speed performance, with nearly 2.4% savings possible from speed optimization compared to 1.5% for altitude optimization. Few barriers exist to some of the strategies such as step climbs and lower speeds, making them attractive in the near term. As barriers are minimized, speed and altitude trajectory enhancements promise to improve the environmental performance of the aviation industry with relative ease.

Acknowledgements

This work was supported by the Office of Environment and Energy, U.S. Federal Aviation Administration, under FAA Cooperative Agreement No. 06-C-NE-MIT, Amendment Nos. 017 and 026.

Disclaimer

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the FAA, NASA, or Transport Canada.

Contents

Acronym a	nd Variable Reference
Chapter 1	Introduction7
1.1 M	lotivation7
1.2 R	esearch Goals
1.3 R	esearch Overview9
Chapter 2	Background12
2.1 A	ircraft Performance
2.1.1	Altitude Effects
2.1.2	Speed Effects
2.2 C	auses of Sub-Optimal Flight Trajectories14
2.2.1	Atmospheric Conditions
2.2.2	Airline and Pilot Planning15
2.2.3	ATC and Airspace Limitations15
2.3 A	ltitude and Speed Profiles Today16
Chapter 3	Analysis Method
3.1 D	ata Sources20
3.1.1	Flight Path Data20
3.1.2	Atmospheric Data
3.1.3	Aircraft Performance Data24
3.2 S	election of Study Cases
3.2.1	Selection of Flights for Analysis
3.2.2	Selection of Aircraft for Analysis28
3.2.3	Final Flight Selection
3.2.4	Levels of Profile Improvement Analyzed
3.3 F	light Performance Analysis Tool
3.3.1	Identification of Cruise Leg33
3.3.2	Initial Weight Estimation34
3.3.3	Calculation of SAR at Any Flight Condition35

3.3.4 Fuel Burn Calculation for Actual Flight Profiles	
3.3.5 Developing the Optimum Speed and Altitude Profile	41
3.3.6 Developing the Step Climb Profiles	44
3.3.7 Developing the LRC Speed Profiles	
3.4 Regional Jet and Turboprop Analysis	46
Chapter 4 Desults	40
4.1 Analysis Tool Output	
4.1 Analysis 1001 Output	
4.2 Individual Results	
4.3 Aggregate Results	
4.4 Regional Jet and Turboprop	
4.4.1 Regional Jet Results	
4.4.2 Turboprop Results	
4.5 System-wide Benefit Potential	73
Chapter 5 Discussion of Results	
5.1 Maximum Improvement Potential	
5.2 Altitude Effects	
5.2.1 Optimum Altitudes	
5.2.2 Step Climbs	
5.3 Speed Effects	81
5.3.1 Optimal Speed	81
5.3.2 LRC Speed	82
Chapter 6 Operational Parriers and Mitigations	94
6.1 Altitude Improvements	
6.1.1 Cruise Climbs	
6.1.2 Stap Climbs	
6.2 Spood Improvements	
6.2.1 Custom Aircraft Speeda	01
6.2.2 Speed Deduction Efforts	٥٥
0.2.2 Speed Reduction Enorts	88
Chapter 7 Conclusion	90

Appendices		
Appendix A.	Aircraft SAR Sensitivity to Speed and Altitude	93
Appendix B.	Aircraft Optimal Altitude Versus Weight	94
Bibliography		

Acronym and Variable Reference

A320	Airbus A320-200
ADS-B	Automatic Dependent Surveillance Broadcast
ATC	Air Traffic Control
B737	Boeing 737-700
B752	Boeing 757-200
BTS	Bureau of Transportation Statistics
CI	Cost Index
C_D	Total Drag Coefficient
C_{D0}	Zero-Lift Drag Coefficient
C_L	Lift Coefficient
c_T	see TSFC
CRJ2	Canadair CRJ-200
D	Drag
DH8D	Bombardier Dash 8 Q400
F	Fuel Burn Rate
FAA	Federal Aviation Administration
ETMS	Enhanced Traffic Management System
FMS	Flight Management System
GHG	Greenhouse Gas
IFR	Instrument Flight Rules
Κ	Induced Drag Factor
L/D	Lift-to-Drag Ratio
LRC	Long Range Cruise
MD82	McDonnell Douglas MD-82
MRC	Maximum Range Cruise
MTOW	Maximum Takeoff Weight
NARR	North American Regional Reanalysis
NAS	National Airspace System
NOAA	National Oceanic and Atmospheric Administration
NOMADS	National Operational Model Archive & Distribution System
q_{∞}	Free-stream Dynamic Pressure
RNP	Required Navigation Performance
RVSM	Reduced Vertical Separation Minima
S	Reference Area (Wing Planform Area)
SAR	Standard Air Range
SFC	Specific Fuel Consumption
SGR	Standard Ground Range
Т	Thrust
TSFC	Thrust Specific Fuel Consumption
V_{∞}	Free-stream Velocity
W	Weight
Θ	Vertical Flight Path Angle (Climb or Descent Angle)

Chapter 1

Introduction

1.1 Motivation

Improving aircraft operational efficiency has recently become a dominant theme in air transportation, as the recent social and political climate has pushed for reduced environmental impact and energy concerns have encouraged decreased reliance on fossil fuels. Mounting scientific evidence of global climate change has spurred increased awareness of the importance of manmade greenhouse gas (GHG) emissions such as *CO*₂, resulting in significant pressure to reduce emissions. While aviation currently contributes 3% of transportation GHGs, this fraction is expected to increase as other transportation modes more easily adopt environmentally sustainable practices (EPA, 2007). Additionally, air transportation is growing at a rapid pace of approximately 5% per year, further adding to the importance of aircraft emissions and corresponding pressure to reduce them (IATA, 2010). Transportation's increasing thirst for fossil fuels has simultaneously generated substantial concerns about the future of the world's energy supply, driving up the cost of petroleum in a trend that is expected to continue (Hirsch, 2005). Environmental concerns have resulted in government pressure to reduce fuel consumption, and increased fuel prices have pushed aircraft operators to find margins for performance improvements.

These factors have resulted in efforts to improve the efficiency of the US air transportation system, which consumed 11.34 billion gallons of fuel in 2009 (BTS, 2010). Efforts to modernize aircraft fleets are limited by the extremely slow and expensive process of new aircraft adoption, which can take decades (Kar, Bonnefoy, & Hansman, 2010). Major infrastructure improvements like the Next Generation Air Transportation System (NextGen) promise efficiency improvements but also face long implementation timelines. Operational improvements, however, remain a viable means of improving environmental performance in the near term.

One operational improvement technique involves increasing the fuel efficiency of flights by improving current cruise flight trajectories. Literature review revealed that most prior work on evaluating operational fuel efficiency has focused on optimization of the descent phase; little has been done to examine the altitude and speed trajectory performance in cruise based on actual flight data. Aircraft performance is tightly linked with airspeed and altitude, so improvements in these dimensions can potentially provide significant increases in efficiency without dramatic changes in infrastructure or routing. Technical and operational barriers will limit the actual success of such measures and must also be considered. A quantification of the potential benefit of various vertical and speed improvement strategies, as well as a discussion of the barriers to their implementation, would provide useful insight into the available improvement potential of such measures, and could help direct near term efforts to increase efficiency. A quantified benefit pool would also help to determine if efforts to improve speed and altitude efficiency are justified.

1.2 Research Goals

This research attempts to accomplish the following goals:

1. Establish an upper bound on the performance benefits attainable in today's airspace system through changes in the cruise speed and altitude trajectories.

- Quantify the potential benefits of various cruise speed and altitude trajectory improvement strategies;
- 3. Identify barriers that may restrict the effectiveness of these strategies.

The purpose of this research is to identify a pool of potential benefits that can be gained from speed and altitude trajectory improvements.

This research primarily attempts to quantify benefits of cruise flight operational improvements to the speed and altitude dimensions. In this research, a benefit is meant to imply a reduction in fuel burn due to a speed or altitude improvement relative to the actual unimproved flight. As CO_2 is directly related to the amount of fuel burned, reduction in fuel consumption implies a reduction in carbon emissions as well. Therefore this analysis answers the question: How much can fuel burn and carbon emissions be reduced in cruise flight if aircraft are operated nearer to or at their optimum speed and altitude?

While an analysis comparing generic trajectories to improved ones would provide some insight, inclusion of actual flight path data is critical in determining how far the system really is from optimal. The key aspect of this research is a detailed comparison between actual flight trajectories and corresponding more efficient trajectories, thus giving the most realistic estimate of improvement potential. Identification of implementation barriers helps establish which optimization techniques are most promising for the near term. These considerations are meant to develop results which are practical and directly applicable to the US air transportation system.

1.3 Research Overview

The techniques available for reducing fuel consumption over any part of the flight can be ultimately separated into two groups: lateral path minimization and aircraft performance improvements. Path minimization simply involves minimizing the distance flown, thus

potentially reducing the amount of time an aircraft is burning fuel. Aircraft performance improvements involve changing the rate of fuel consumption during the flight, which can be achieved via airframe modifications like winglets, or by operating an aircraft nearer to its ideal flight condition. While the effects of airframe modifications and lateral improvements are fairly well understood by manufacturers and flight planners, little has been done to examine how far aircraft are actually operated from their ideal flight condition. Air density (altitude) and aircraft Mach number (speed) have significant and correlating effects on aircraft performance, and thus were simultaneously chosen as dimensions to examine.

To provide an assessment of the potential fuel efficiency benefits from improved speed and altitude a baseline of current operational performance was created. Flight and atmospheric data were combined with aircraft performance tools to estimate fuel burn on a per flight basis. Archived data from the FAA's Enhanced Traffic Management System (ETMS) was used to provide flight path information. Atmospheric data were acquired via the National Oceanic and Atmospheric Administration's (NOAA's) atmospheric operational model. Lissys Piano-X, a professional aircraft analysis tool, was utilized to characterize the fuel burn performance of various aircraft. A custom analysis code was then developed to provide fuel burn estimates for each point in the flight path, integrating to find the fuel burn in the flight's cruise phase. Performance data was used to then develop theoretical altitude- and/or speed-optimum flight paths, whose fuel burn was also calculated and compared with the original fuel burn. These comparisons serve as the basis for analysis results. More details about this analysis process are discussed in the methodology.

The vertical and speed profiles of commercial aircraft in cruise flight were the subject of evaluation in this research. These profiles were analyzed alongside proposed improved profiles. The analysis exclusively examined the cruise phase of flight, ignoring the climb and descent at the beginning and end of each flight. While the bulk of the effort involved development and

implementation of a robust tool to analyze flights, the latter work focused on the results of this analysis, looking for trends and evaluating the difference in impact across various improvement levels. Trends were examined by comparing results across categories like stage length, airline, and aircraft type. The primary evaluation metric used to compare analysis results was simply the percentage reduction in cruise fuel burn possible, from the actual to the improved flight trajectory. Comparisons were drawn between the various types of trajectory improvements in an attempt to identify the best performing candidates.

Chapter 2

Background

2.1 Aircraft Performance

This work focuses on the potential to increase aircraft performance via changes in flight operations. In this context, increased performance is meant to denote a decrease in cruise fuel burn, and correspondingly a decrease in CO_2 emissions, from what is commonly observed in today's National Airspace System (NAS). Various means exist which to help improve aircraft performance. One strategy of performance improvement involves modifying the aircraft *itself*. Aerodynamic modifications, powerplant upgrades, and weight reduction efforts are a few examples of improvement methods which utilize altering the aircraft's physical characteristics. Alternatively, the *way* that the aircraft is flown can yield performance improvements without requiring any changes to the vehicle itself. These operational strategies can potentially generate environmental and financial benefits for all flights in the NAS. Research also suggests that operational improvements are more likely to have an impact in the near term when compared to aircraft modifications, which are slow to migrate into the system (Kar, Bonnefoy, & Hansman, 2010).

2.1.1 Altitude Effects

An understanding of how operations affect performance is required before improved operations can be suggested. The first of two flight dimensions being considered for adjustment is altitude. Aircraft are designed to operate optimally at a specific altitude that is dependent on weight. As altitude increases, the air density decreases, and it is this density which is the underlying parameter affecting altitude performance. Both the aerodynamic and engine performance are affected by altitude, however only aerodynamic performance is influenced by the weight of the aircraft. The lift coefficient, C_L , generally has a single constant value at which the lift-to-drag ratio of the aircraft is optimized, and is defined as:

$$C_L = \frac{L}{q_{\infty}S}$$

L is the lift of the aircraft, equal to the weight in steady level flight, *S* is a constant reference area, and q_{∞} is the freestream dynamic pressure defined by:

$$q_{\infty} = \frac{1}{2} \rho V_{\infty}^{2}$$

Here, ρ is the air density, and V_{∞} is the freestream velocity, or airspeed. As fuel is burned and weight decreases during cruise, dynamic pressure must also be reduced if the lift coefficient is to be maintained at its optimum value. Therefore, for a given cruise velocity, the optimum density decreases throughout the flight, corresponding to an increase in altitude. This result offers a basic understanding of what an ideal altitude profile would look like: a slow and steady climb, or cruise climb. Even when the optimal altitude might not change much during cruise, this relationship still makes clear that the ideal cruise altitude for a certain aircraft can vary significantly depending on the amount of fuel, passengers, and cargo onboard.

2.1.2 Speed Effects

The second dimension under consideration is speed. As evidenced by the definition of dynamic pressure, the airspeed is closely linked with altitude in part of its effect on the aerodynamic performance. In addition to the effect of speed on the optimum lift-to-drag ratio, speed has several more complex relationships with aircraft performance. As most commercial jets operate in the high subsonic realm, relatively minor speed increases can often result in drastic increases in drag, due to compressibility effects and shocks created in transonic flow. In addition, turbofan engine performance is also very sensitive to Mach number. Because optimal speed is not sensitive to weight, it remains constant in zero wind environment. However, aircraft performance for any given flight is a function of fuel burned over ground distance flown, so winds do have an effect on optimal airspeed: headwinds increase the optimal speed, while tailwinds decrease it.

2.2 Causes of Sub-Optimal Flight Trajectories

If a simple change in flying behavior can provide cost savings and improved environmental performance, why are sub-optimal trajectories being flown today? A multitude of reasons exist, but they can ultimately be grouped into atmospheric conditions, airline and pilot planning, and ATC and airspace limitations.

2.2.1 Atmospheric Conditions

One of the most direct and probably most obvious reasons for diverting from optimal, atmospheric conditions can reroute aircraft for safety reasons or simply for comfort. While commercial aircraft fly above most weather, severe cells can extend to dramatic heights and require altitude diversion. Weather problems at any level can also cause delays, which often result in pilots flying faster than ideal in an attempt to make up time. Even more common, however, is turbulence avoidance. When aircraft encounter persisting pockets of rough air at the given cruise level, pilots will often request a lower or higher flight level in search of smooth

air for increased cabin comfort. While this diversion often comes at the expense of increased fuel burn, passenger comfort is generally considered worth any minor decrease in performance. Given the importance of safety and comfort, atmospheric disruptions are unlikely to be avoided.

2.2.2 Airline and Pilot Planning

While ATC exists to maintain safe and efficient traffic flow, it still serves the role of a guide, and most operating decisions are left to the pilots and their airlines. One factor that often drives cruise speeds above the best economy (fuel-optimal) speed is the cost index (CI). CI's exist as a means to represent the balance between the cost of fuel and the cost of time on a given flight. Flying faster than best economy will cost extra fuel, but the flight time will be reduced, resulting in a reduction in labor and maintenance costs, as well as increasing the availability of the aircraft for other operations. Airlines seek to reduce costs, which means flying at the speed dictated by the CI and not only by the speed that reduces fuel and emissions. This financecentric operating policy is partly responsible for operating excursions beyond the performanceoptimal speed profile.

Another issue possibly resulting in higher cruise speeds is timeliness. When congestion or other issues cause delays, pilots often add speed in order to ensure connections are made or simply to improve passenger satisfaction with the airline.

Finally, pilots are balancing many operational considerations in speed and altitude decisions and may not fully appreciate the fuel efficiency impact of these decisions. Airlines need to establish clear policies and plans that encourage optimal flight paths. Without this planning, pilots may not be aware of the potential for increased performance and will not take action to seek it out.

2.2.3 ATC and Airspace Limitations

The constraints of the NAS play a major role in preventing ideal trajectories from being flown. The lateral depiction of aircraft on controllers' radar screens limits the vertical

situational awareness. In addition, the task of ensuring safe separation becomes difficult when aircraft are free to move in all three dimensions so aircraft in vertical transition result in higher controller workload. In an attempt to manage complexity for controllers, aircraft are generally limited to level flight paths on common airways and altitude levels, with special caution and attention given to the occasional climbs and descents between flight levels (Histon & Hansman, 2008). These vertical separation limits have been reduced to 1,000 ft with implementation of the Reduced Vertical Separation Minima (RVSM) in the US. Traffic is separated by direction on alternating flight levels, so the vertical distance between two valid flight levels in a given direction is 2,000 ft under RVSM. As a result, an aircraft attempting to follow an ideal altitude trajectory in today's airspace with have to undergo 2,000 ft step climbs.

ATC organizes aircraft along known routes to simplify handling of traffic. This poses problems in congested airspace and when the speeds of aircraft along the route do not match. As a faster aircraft approaches a slower one from behind, one aircraft is forced to change speeds, divert laterally, or divert vertically. In the case of changed speeds or vertical diversion, the aircraft speed or altitude profile may diverge from optimal and the performance suffers. In very congested airspaces, a given flight level may be "full" of other traffic, potentially affording no possibility of certain aircraft to operate near their ideal altitude.

Another potential for reduced performance results from the priorities of controllers. Controllers exist first to ensure safety, and then to improve capacity. A focus on throughput could potentially result in aircraft flying at undesirable speeds or flight levels, and may be responsible for some of the reduced performance.

2.3 Altitude and Speed Profiles Today

The issues plaguing cruise altitude and speed profiles can best be described with representative data. Figure 1 shows the cruise segment of a flight used in this analysis, a Boeing 757-200 traveling from Boston to San Francisco. On the left is the altitude profile, and on the

right, the calculated airspeed in Mach. The blue lines represent the actual path flown, while the green represents the optimal path. This flight represents of the forces at play. In due diligence, the pilot makes a 2,000 step up at about 1,200 nm into the flight, the smallest increment possible under RVSM. This step is in line with staying as close to optimal as is possible. Unfortunately, at approximately 1,500 nm, the pilot steps down again. Clearly this was not a performance enhancing choice—most likely this diversion was caused by ATC due to a traffic conflict, or by turbulence at the higher flight level that the pilot hoped to escape by returning to the last known smooth flight level. The speed profile illustrates a trend all too common in today's cruise operations: flying fast. Despite some noise in the data, it is clear the aircraft was traveling at approximately Mach 0.80 when the best economy speed for the 757 is approximately Mach 0.76.



Figure 1. Actual and ideal flight profiles for a Boeing 757-200 from Boston to San Francisco.

Another sample cruise flight profile is shown in Figure 2. This flight represents a shorter trip of a Boeing 737-700 from Los Angeles to Chicago. In this altitude profile, the aircraft remained at FL390 for the entirety of the flight. As fuel was burned, however, the ideal altitude rose to over FL400. Whether or not the pilot intended to seek the optimal altitude condition is not known, because step climbs of 1,000 ft are not normally given under current RVSM

restrictions. Here, the flight was likely limited by these ATC constraints to a level flight profile. The speed profile is similar to the prior case—the aircraft flies faster than best economy for the entirety of the flight. The cause of the faster speed could be due to CI selection, or simply pilots attempt to make up time.



Figure 2. Actual and ideal flight profiles for a Boeing 737-700 from Los Angeles to Chicago.

A third example is depicted in Figure 3 for a MD-80 between New York and Chicago. This flight represents a common theme seen on shorter, busy routes. The difference between the initial cruise altitude and the final optimal altitude is only 1,000 ft, so a step up was not feasible. Also, the pilot prematurely stepped down 2,000 ft nearly 100 nm before starting the descent from cruise. The early step down was seen on many other flights traveling this route, and is likely due to ATC routing procedures into the busy Chicago terminal area.



Figure 3. Actual and ideal flight profiles for a McDonnell Douglas MD82 from New York to Chicago.

Chapter 3

Analysis Method

3.1 Data Sources

Data from several sources were collected to create the necessary basis for the analysis. As actual flights were being analyzed, detailed flight track information was required. Atmospheric data was needed to completely characterize each flight. Finally, comprehensive aircraft performance data was needed to accurately assess the aircraft's performance on each flight.

3.1.1 Flight Path Data

The FAA monitors and records the entire NAS traffic picture, including the position and altitude of each aircraft in flight. This system, the Enhanced Traffic Management System (ETMS), is the tool used by ATC to manage the flow of traffic in the NAS (FAA, 2009). This information is made public to certain organizations for research related purposes, and a select sample of this data was made available for use in this effort. The dataset includes flight path and flight plan information for one entire day of US flights, including those originating or terminating outside the country. The data covers the entire 24 hour period on September 21, 2009, and includes nearly 53,000 individual flights.

The ETMS data is most helpful in that it provides location, altitude, and time values for each flight at approximately one minute intervals. However, as the location data is gathered from primary radar returns gathered from various radar stations, the location accuracy is limited.

Latitude and longitude accuracy was reported to the nearest minute, a length equal to a nautical mile in latitude. The aircraft analyzed travel at approximately seven nautical miles per minute, so some of the distance calculations between points (and corresponding ground speed calculations) could easily yield errors of 10% or more. However, calculations made over the entire length of the flight meant that these errors would average out and were not cumulative, yielding much more accurate average values. The implication of this fact is that ground speed calculations generated significant short term variation, but smoothing these spikes over larger time frames produced appropriate results. An example of the raw and processed groundspeed data is shown in Figure 4.



Figure 4. Example of a common speed profile spanning the entire flight, showing the unprocessed data (blue) and the smooth time-averaged data (red).

Ultimately, the ETMS flight path information successfully provides distance, groundspeed, and altitude information for all US domestic flights operating under ATC command, namely instrument flight rules (IFR) traffic.

3.1.2 Atmospheric Data

While groundspeed is important in determining the speed of an aircraft, it does not provide enough information to completely determine an aircraft's velocity through the air. Without detailed onboard flight data, airspeed must be calculated with the combination of groundspeed and wind information. Therefore, wind speed and direction data were needed for locations along the entire flight path, at various altitudes. In addition, temperature of the ambient are was needed to calculate the Mach number, which is more pertinent to transonic aircraft performance than just airspeed alone.

Wind and temperature data was gathered from the NOAA, which maintains a detailed US atmospheric model. Data was accessed through the NOAA's web-accessible model data archive, called the National Operational Model Archive and Distribution System (NOMADS). This system was created to address a growing need for data from NOAA's numerical weather prediction and climate models (NOAA, About NOMADS, 2010). The model accessed was the North American Regional Reanalysis (NARR), which assimilates a large amount of observational data to create a detailed nationwide picture of past atmospheric conditions. The model includes observations from radiosondes (instruments on weather balloons), surface sensors, dropsondes (instruments dropped from aircraft), aircraft sensors, and satellites (NOAA, NCEP North American Regional Reanalysis , 2010). Because this model does not make weather predictions, and includes the most complete set of observations available, it was selected as the best source for wind and temperature data.

NARR data is provided in the GRIB format, a standard for large sets of meteorological data, and is organized on a Lambert conformal conic projection type grid. NARR provides data across the entire US in a lateral spacing of 32 km, at 3-hour time intervals, and at variable altitude spacing. The vertical spacing of data is the finest at altitudes less than 10,000 ft and greater than 30,000 ft. This trend suggests that higher frequency data was favored near the ground and at the altitudes flown my commercial aircraft, where it is likely most often used. The vertical spacing of available data versus the altitude is shown in Figure 5. Each blue dot here represents a pressure altitude (along the x-axis) where data exists. The y-value represents the vertical

distance to the next data point. The implication here is that wind and temperature data is spaced at approximately 2,000 ft vertically for most flights, which occur between 30,000 and 40,000 ft.



Figure 5. Vertical data spacing as a function of absolute altitude. Each point represents an altitude where NARR data is available.

The implication of the Lambert conformal grid is that the spacing between data points is approximately constant in distance on the surface in the US, but longitude spacing is variable because longitude lines become closer as distance increases from the equator. This grid is most appropriate for the US, but appears distorted when plotted using the common Mercator projection, where longitude and latitude lines form a perpendicular grid. An example of this weather information is shown in Figure 6, which shows the surface temperature. The warm temperatures on land help distinguish the southern US coastline. The stretched nature of the grid is due to points of longitude being farther apart at more northern latitudes. Wind component data was gathered along with temperature in this format, and at various altitudes and times.



Figure 6. Surface temperature across the US. This plot provides an example of the scope of data provided by NARR.

3.1.3 Aircraft Performance Data

No analysis of altitude or speed sensitivity could take place without understanding an aircraft's detailed performance at various conditions. Vehicle performance was characterized using the tool Lissys Piano-X, a professional aircraft analysis tool capable of providing detailed performance information for a variety of commercial aircraft flying entire mission profiles or at single point conditions. Piano-X is not able to perform calculations on customized profiles like the ones being analyzed in this research, so a custom flight profile analysis tool was developed separately. Still, Piano-X does contain detailed performance information for various aircraft at any given weight, altitude, and speed, which was sought in this analysis. At any such steady condition, Piano-X can provide fuel burn rate, thrust required, specific fuel consumption (TSFC), lift coefficient (C_L), drag coefficient (C_D) and its components, lift to drag ratio (L/D), and standard air range (SAR).

The measure of performance most critical to analyzing fuel burn over a given distance is SAR, which is expressed in distance traveled (nautical miles) per mass of fuel burned (kilograms). This metric gives a direct relationship between fuel burn and distance, which is useful in integrating the total fuel burn over a given flight path. Because SAR is only referenced to the distance flown through the air, it does not alone capture the vehicle's efficiency over the ground. The wind data was later incorporated at each point in flight to correct SAR from air distance flown to ground distance flown, giving rise to a similar instantaneous metric called standard ground range (SGR). To take into account fuel burn changes in climbs and descents, more information was needed relating engine thrust to fuel burn. Therefore, the aircraft's thrust specific fuel consumption (TSFC) was also retrieved from Piano-X along with SAR.

Throughout the flight, three main performance-related dimensions are in constant flux: weight, speed, and altitude. Under ideal circumstances, the SAR and TSFC data would be characterized across all these three dimensions. Unfortunately, while Piano-X does have an interface for providing these inputs, it can only accept inputs manually and must be run one case at a time via a graphic application interface. This makes a complete factorial data collection effort across all three dimensions extremely impractical. Instead, data was collected across the speed and altitude dimensions at a constant weight, and a separate calculation was made to characterize sensitivity to changes in weight. Speed and altitude were chosen as the primary dimensions for data collection because their effect on performance is the focus of the analysis, and the effect of weight is better understood.

Before running sweeps of altitude and speed, the ideal altitude and speed for the given weight were determined. The weight selected for the tables was the nominal weight input chosen by Piano-X for each aircraft. This weight represented a mid-range point for most flight operations, and was 80-90% of maximum takeoff weight (MTOW). The detailed mission profile mode was then run and the optimum altitude was observed. Nominal Mach numbers were found by observing Piano-X's cruise Mach selections for economy and long range cruise (LRC) modes. Based on these nominal speed and altitude selections for the given weight, a grid of

speed and altitude points were established around this point. Altitudes were spaced at 2,000 ft intervals, and Mach numbers were spaced at hundredths of a Mach. The upper and lower bounds for each of these dimensions were selected by observing the typical range of speeds and altitudes flown by each aircraft in Piano-X, and then applying some margin to extend the grid further. This setup allowed for nominal cruise conditions and deviations to be handled with ease, although inevitably not encompassing all possibilities. The resulting tables contained 7-8 altitude conditions and 10-12 speed conditions, yielding close to 100 individual cases to be run and recorded for each aircraft. Each SAR table was then normalized by the highest value such that only the relative change from maximum was recorded at each point. The altitude and speed lookups were also likewise normalized to zero at the ideal condition. This step improved readability of the tables and made variability easy to observe. The maximum SAR value was stored separately so that the table values could be converted back to absolutes easily. Contour plots showing these SAR data are shown in Appendix A. An example of one of these plots is shown in Figure 7. Each solid line represents a 1% change in SAR.



Figure 7. SAR contour plot for Aircraft 3. Starting at the optimum SAR (marked by a single point) and moving outwards, each line represents a 1% decrease in SAR from optimal.

A weight-versus-altitude correlation was then developed. This was achieved by programming the design range and payload into Piano-X, and then providing a wide range of

optional altitudes for the aircraft to fly at. Piano-X optimizes altitude automatically in the mission profile mode, and the resulting altitude steps at various weights were recorded to develop a relationship between the two. Weight was recorded at the beginning of each new step altitude, after the aircraft had finished climbing. As expected, the ideal altitude is approximately linear with weight. Plots showing the resulting data and linear fits are shown in Appendix B. An example for Aircraft 3 is shown in Figure 8.



Figure 8. Weight versus ideal altitude estimates from Piano-X for Aircraft 3.

Finally, the induced drag factor (K) was determined for each aircraft. This factor is used in the standard drag polar equation as a scalar on the square of C_L to indicate the influence of lift on an aircraft's drag. K was estimated by thus dividing the lift-induced drag component (C_{Di}) by the square of C_L (both are outputs of Piano-X) for a variety of flight conditions. The need for this variable is related to characterizing the effect of weight on SAR, which is described later.

3.2 Selection of Study Cases

Given the available resources, the potential scope of the analysis was vast. ETMS data made thousands of flight paths available analysis, and Piano-X contained performance data for nearly any commercial aircraft. However, in order to keep the scope of analysis within reasonable limits, only a fraction of these flights and aircraft was selected for analysis. Selection of aircraft and flights for the purposes of this research was based on two overarching goals: the analysis should be applicable to the largest possible group of current domestic flights, and the cases should provide enough diversity such that meaningful differences in the results of various aircraft and route types could be discernable. A tool was developed to allow easy searching of the ETMS database, given inputs of departure/arrival airports and aircraft type. This tool allowed scouring of the available flight paths to identify popular routes and common aircraft types, and was instrumental in the selection of the cases.

3.2.1 Selection of Flights for Analysis

Flights were initially selected by city pairs. Popular city pairs were chosen in line with the goal of capturing large portion of the US traffic, and thus providing results which would provide the most relevant improvement potential. Diversity in route stage length was also sought along with route flight volume. This diversity allowed differences in the improvement potential across various stage lengths to be made apparent. Additionally, certain aircraft types are often tied to certain routes, so a mixture of route lengths increased the types of aircraft available for consideration. An important fact to note about this analysis is that the pairing of aircraft type with each individual flight was maintained. That is, the aircraft type operated on a given flight was used in the performance analysis of that flight. No type substitutions were made, thus increasing the authenticity of the results.

An array of popular city pairs was identified that represented a range of stage lengths and a variety of aircraft. This selection contained 14 different city pairs and 1234 total flights. From this, a set of popular aircraft was identified, and then the set of flights was narrowed down to those operated by the selected aircraft.

3.2.2 Selection of Aircraft for Analysis

The selection of aircraft types is closely linked with the flight selection, since each individual flight corresponds to one type of aircraft. In an attempt to capture as many flights from the

initial selection as possible, the aircraft type demographic was characterized. Figure 9 shows the aircraft type breakdown for the initial flight selection.



Figure 9. Total occurrences of all aircraft types present in the initial selection of flights

The Boeing 737 is the most produced jet airliner in history, and the most common in this subset, making it an obvious choice (Kingsley-Jones, 2009). The 737-300 (B733), -400 (B734), -500 (B735), -700 (B737), and -800 (B738) are all variants of the Boeing 737 in this subset, however the -700 model was specifically chosen for the analysis due to its prevalence in the data. The 737's chief competitor, the Airbus A320 (A320), also appeared frequently and was likewise chosen. These aircraft represent a large portion of the single-aisle airliner market and are ideally sized to operate on many US routes, making them very common. The Boeing 757-200 (B752), a larger capacity single-aisle aircraft, sees significant utilization on US domestic routes, and was selected for analysis. The McDonnell Douglas MD-82 (MD82) was also chosen because it represents an older aircraft that is still in frequent use. Inclusion of this age diversity allows an extra dimension over which to compare the final analysis results. The Canadair CRJ-

200 (CRJ2) and Bombardier Dash 8 Q400 (DH8D) were chosen as popular representatives for regional jets and regional turboprops, respectively. Table 1 lists the aircraft selected for analysis.

Selected Aircraft
Boeing 737-700
Airbus A320
Boeing 757-200
McDonnell Douglas MD82
Canadair CRJ-200
Bombardier Dash 8 Q400

Table 1. List of selected aircraft for analysis.

3.2.3 Final Flight Selection

Identification of the aircraft for analysis helped to narrow down the flight cases, as only the flights flown by the selected aircraft could be used in the analysis. Additionally, some route-aircraft combinations overlapped in type and stage length and were removed to make the dataset more manageable. Other flights contained corrupted data and could not be used; these were omitted as well. The CRJ and Dash 8 aircraft were not included in the standard analysis and were handled separately; this process is discussed in detail later. The final city pair and aircraft type combinations actually used in the analysis are shown in Table 2. This flight dataset includes a total of 257 domestic US flights which are used as a basis for all calculations in the analysis.

City Pairs	Stage Length (nm)	B737	A320	B752	MD82
Atlanta - Miami	517	5		22	
Washington DC - Chicago	530		19		
New York - Chicago	641	14	30		33
Washington DC - Dallas	1033				25
Los Angeles - Chicago	1512	12	11	18	
New York - Los Angeles	2144	6	26	28	
Boston - San Francisco	2343			8	

Table 2. Final aircraft and city pair cases chosen for analysis.

3.2.4 Levels of Profile Improvement Analyzed

Selection of the actual flights to be used in the analysis provided a clear subset of input data to be processed. Equally important were the types of analysis to be run on each flight. The overarching goal of each analysis case was to determine the fuel burn, which was used for comparison. The first case was simply an analysis of the existing flight path in its unaltered form. This result served as the baseline reference to which all subsequent modifications were measured, and represents the state of cruise operations today. Next, the optimal speed and optimal altitude case was selected to provide an upper bound on the possible cruise improvement.

Other combinations of altitude and speed improvement were chosen to provide insight into the potential of various operational mechanisms. Because of the interactions between speed and altitude, the benefit achieved by optimizing both parameters simultaneously is not equal to the sum of the benefits of optimizing each separately. By analyzing them separately, the interaction effects could be identified, and the influence of each dimension's optimization could readily be observed. Therefore, an optimal speed but unaltered case was chosen, as well as an unaltered speed but optimal altitude case.

Still, all of these cases represent either the best or worst cases for altitude or speed. These types of operations would likely be difficult to implement, so intermediate cases were needed. Step climbs are operational procedures that represent moderate attempts to improve altitude optimality. Today, RVSM allows at best 2,000 ft step climbs. Therefore, an altitude improvement case allowing 2,000 ft step climbs was selected. To help examine the sensitivity of step climb benefits to step size, a finer 1,000 ft step climb case was chosen as well. The smaller step case also represents a more realistic improvement in altitude flexibility than a pure cruise climb. For example, if RVSM were available on one-way routes, then vertical separation

in the given direction would indeed be 1,000 ft. For these step altitude cases, speed was left unaltered in order to isolate the step climb benefit.

Flying fuel-optimal Mach generally means flying at the maximum range cruise (MRC) setting. However, this is an unlikely choice by most operators, who seek to minimize costs. Operators select speed using a cost index (CI), which is the ratio of the cost of time to the cost of fuel. Before each flight, CI is programmed into the FMS which then selects a speed to minimize cost. This speed is inevitably faster than MRC, as time always has a cost. A more likely speed selection is long range cruise (LRC) setting, which trades off approximately 1% of fuel efficiency for 3-5% increase in speed. The LRC Mach represents a time-sensitive CI, even corresponding to a CI value higher than most operators are likely to normally select (Roberson, Root, & Adams, 2007). LRC contrasts MRC well by representing typical scenarios, such as recovery time lost in delays, where operators are pressed for time and hurrying. LRC is also a programmable setting common to all aircraft, making it a good choice for comparison across various aircraft types. For these reasons, a LRC speed case was created. The hypothesis is that many aircraft still fly faster than this speed, making LRC an improvement from typical performance. Unlike the more drastic fuel-optimal speed option, LRC provides a more conservative efficiency improvement case for analysis. The unaltered original altitude profiles were paired with the LRC speed case for consistency and to isolate the LRC benefit. The resulting speed and altitude analysis cases are shown in Table 3.

Case	Speed	Altitude
1	Optimal	Optimal
2	Optimal	Unaltered
3	Unaltered	Optimal
4	Unaltered	Step 1000 ft
5	Unaltered	Step 2000 ft
6	LRC	Unaltered

Table 3. Speed and altitude cases chosen for analysis.

3.3 Flight Performance Analysis Tool

With all the necessary data in hand, the next step was to develop a method to procedurally step through each flight profile, calculate the instantaneous fuel burn at each point, and integrate the total fuel burn over the flight. Next, the original flight path was modified by improving one or both of the speed and altitude dimensions, and then its fuel burn calculated. The final step was simply comparing the fuel burn of the improved trajectory to the original and identifying fuel burn savings.

3.3.1 Identification of Cruise Leg

Most flight profiles contain data starting when the aircraft climbs above 2,000-5,000 ft, and ending when the aircraft descends below a similar altitude. Because this analysis was limited to the cruise leg only, the cruise data had to be separated from the climb and descent legs. An automated means of finding these start and end points was ineffective because aircraft often level off at different parts of the climbs and descents. Instead, the start and end points were determined subjectively, by viewing the entire altitude profile and observing where the long plateau-like phase gave way to steadily sloping climbs and descents. The cruise phase was manually selected in this manner for each flight individually, to ensure no other parts of the flight were accidentally included and thereby contaminating the results. Figure 10 gives an example of a typical altitude profile, with lines denoting the chosen start and end of cruise. The data between the lines was used in the analysis.



Figure 10. Typical altitude profile, with vertical lines marking the start and end of cruise.

3.3.2 Initial Weight Estimation

Before any performance calculation could be calculated, the weight of the aircraft needed be known. This data was not available from ETMS, and detailed information from the FMS or dispatchers was unavailable, as previously mentioned. The Bureau of Transportation Statistics (BTS) provides monthly aggregate data regarding aircraft payload weight; however aggregate data was not appropriate in this analysis because it does not capture the performance-sensitive variation that exists on a per-aircraft and per-flight basis.

In lieu of detailed weight information from the FMS or airline dispatchers, weight was estimated for each flight. A rough assumption was made that the filed cruise altitude was chosen by the dispatchers to correspond with the optimum condition at the start of cruise. In other words, the aircraft dispatchers, who have access to the weight information, ideally selected an initial cruise altitude close to the aircraft's optimal altitude at that weight. Based on this assumption, the weight was calculated using the filed altitude as an input to the linear weightaltitude relationship previously established (and shown in Appendix B) for each aircraft.

This assumption may be incorrect at times as other operational factors may dictate that filed altitudes are selected for a number of reasons unrelated to weight, including weather, airspace, and expected clearances. While the proposed weight assumption is not by any means perfect, it

was the best method that could be determined with the available data. It provides a reasonable estimate of weight for each flight and its corresponding improved variants. Furthermore, the absolute weight value is not critical, because improvements in performance for a given flight are measured across profiles that utilize this same initial weight estimate.

3.3.3 Calculation of SAR at Any Flight Condition

The calculation of SAR for a given flight condition was critical to the success of this analysis. As previously described, SAR is the Standard Air Range, which is ultimately the airplane equivalent of miles per gallon (MPG). SAR is defined as the instantaneous ratio between distance flown through the air and the amount of fuel burned in that distance. When corrected for winds to reference ground distance flown instead of air distance, it is referred to as Standard Ground Range (SGR). This metric can be used to monitor or predict the efficiency of an aircraft at any point in the flight. The integration of its inverse over a given distance yields the total fuel burn over that distance. Therefore it is used in the performance analysis to both monitor efficiency and calculate fuel burn.

The procedure for calculating SAR is also used both in analysis of the actual flight profile and the modified ones, so its mechanics are discussed first here. The SAR tables obtained from Piano-X provide altitude and speed dependent data for one weight; more complex calculation is required to correct those values for arbitrary weight inputs.

The source of the aircraft's fuel burn can be traced to the engines. Their fuel burn is related to the TSFC, which is related to the flight condition (primarily speed and altitude), and the thrust being developed, which is equal to the drag of the aircraft in steady level flight. At a given Mach condition (thus eliminating effects caused by compressibility and transonic changes) the drag polar can simply be reduced do a zero-lift component (C_{D0}) and a lift-induced component based on C_L as follows:

$$C_D = C_{D0} + K C_L^2$$
At a certain altitude and speed, the engine inlet conditions are unaffected by weight, making TSFC nearly constant with weight. C_{D0} is based on the drag of the airframe unrelated to any lift, and thus is also unaffected by weight changes. The only factor influenced by weight is the lift coefficient, C_L , which affects the lift-induced drag component. With these assumptions in mind, the weight correlation was developed.

The process for weight correction involved first calculating C_{D0} using the weight and maximum SAR point used to create the tables, holding this constant, and then using it to recalculate SAR at the new weight condition. This process can be better understood when the various performance dependencies are traced. SAR, which represents distance flown per fuel burned, can be also represented as follows:

$$SAR = \frac{nm}{kg} = \frac{nm/hr}{kg/hr} = \frac{V_{\infty}}{F}$$

Here, *V* is the velocity and *F* is the fuel burn rate. Because velocity is not changing in the correction for weight, the focus moves to the fuel burn rate, which can be expressed by thrust (*T*) and TSFC (c_T):

$$F = T \cdot c_T$$

Since thrust is equal to drag in steady level flight,

$$F = D \cdot c_T$$

Expressing drag in terms of drag coefficient,

$$F = C_D q_{\infty} S c_T$$

Expanding the drag coefficient,

$$F = (C_{D0} + KC_L^2)q_{\infty}S c_T$$

Expanding the lift coefficient in the lift-induced drag component, and knowing that lift equals weight in steady level flight,

$$F = \left[C_{D0} + K \left(\frac{W}{q_{\infty} S} \right)^2 \right] q_{\infty} S c_{T}$$

Finally, C_{D0} and S, the reference wing area (constant for each aircraft), can be simultaneously separated as a constant drag area:

$$C_{D0}S = \frac{F}{q_{\infty}c_{T}} - \frac{K}{S} \left(\frac{W}{q_{\infty}}\right)^{2}$$

Here, *K* was pre-calculated for each aircraft from Piano-X and *S* is a known reference area. This drag area was calculated at the known weight used to generate the SAR table, at the speed and altitude condition of the reference (maximum) SAR point on the table. Again, assuming this drag area is constant, it can be used in the calculation of SAR at the new weight. SAR can be corrected by knowing the drag at the table weight condition and the drag at the new weight:

$$SAR_0 = \frac{V_{\infty}}{F} = \frac{V_{\infty}}{D_{W0} \cdot c_T}$$

$$SAR_1 = SAR_0 \frac{D_0}{D_1}$$

Here, 0 and 1 are used to denote the reference table weight and the new arbitrary weight, respectively. The drag terms are expanded using the previous substitutions:

$$D = \left[C_{D0} + K\left(\frac{W}{q_{\infty}S}\right)^{2}\right]q_{\infty}S = C_{D0}S q_{\infty} + \frac{K}{S}\frac{W^{2}}{q_{\infty}}$$

Finally, the resulting weight correction, assuming constant drag area and *K*, yields:

$$SAR_{1} = SAR_{0} \cdot \frac{\left(C_{D0}S q_{\infty_{0}} + \frac{K}{S} \frac{W_{0}^{2}}{q_{\infty_{0}}}\right)}{\left(C_{D0}S q_{\infty_{1}} + \frac{K}{S} \frac{W_{1}^{2}}{q_{\infty_{1}}}\right)}$$

To further clarify, this correction does not yield the final SAR at the new flight condition. The 0 state is at the weight used to generate the SAR tables, and at the speed and altitude of the ideal condition (max SAR) at that weight. The 1 state is at the new weight, and at the ideal speed and altitude at that new weight. This correction essentially converts the reference (maximum) SAR for the original table weight to the new reference SAR for the new weight. To this point, no inclusion of off-optimum speed or altitude has been introduced, only the reference max SAR points for each speed/altitude table have been used. The off-optimum speed and altitude penalties were then applied after the max SAR was corrected for weight because that ordered approach provided the most accurate results when referenced with cases from Piano-X.

The altitudes used to generate the SAR table were input as pressure altitudes, assuming a standard atmosphere. However, aircraft performance is much more closely linked to density altitude, so an array of corresponding air densities was tabulated alongside the altitude inputs and used for calculations of off-optimal deviations. For a given input into the SAR correction function, the density of the air was computed based on the altitude and temperature. The ideal density (corresponding to the ideal altitude) was also selected based on the weight, and the difference between the two computed. Likewise, the difference is computed between the actual and the ideal Mach. The ideal Mach is selected as that which yields maximum SAR in the table for each aircraft, and was assumed and shown to be nearly constant with weight. Using these density and speed deviations as lookup values, the SAR penalty was found via 2D interpolation of the normalized SAR tables. This penalty was applied to the weight-corrected optimal SAR from above, finally resulting in the SAR for level flight at the input conditions.

The last calculation involved a correction for climbs and descents. Aircraft obviously burn more fuel in climbs, and less in descents. This effect was easily modeled as an increase or decrease in thrust required. Rearranging a previous expression for SAR:

$$T_{lev} = \frac{V_{\infty}}{SAR_{lev} \cdot c_{T}}$$

This calculation yields the thrust in level flight, based on known velocity and specific fuel consumption. The corrected SAR is simply a rearrangement of this equation, but with an additional thrust required component due to the weight of the aircraft:

$$SAR = \frac{V_{\infty}}{(T_{lev} + W \sin \Theta) \cdot c_T}$$

This equation shows that a positive flight path angle (Θ) results in extra thrust required and thus lower SAR, which was expected. Substituting, the final corrected SAR is:

$$SAR = \frac{V_{\infty}}{\left(\frac{V_{\infty}}{SAR_{lev}} + W \sin \Theta c_T\right)}$$

As expected, when the flight path angle is zero, the corrected SAR is equal to the level flight SAR.

The procedure described in this section was in utilized in each segment of every flight. The means of correcting for weight was validated against Piano-X SAR data at selected speed, altitude, and weight points at the bounds of what was expected to be reached in the analysis, with an average error of 1.2%. The worst case scenario was the lowest C_L condition, caused by high Mach, low altitude, and low weight; this weight correction resulted in a 3.4% SAR error compared to Piano-X. Given the low relative error, and considering that all flight calculations used this same process (thus preventing any biased comparison), this method of SAR calculation was deemed acceptable for the purpose of this analysis.

3.3.4 Fuel Burn Calculation for Actual Flight Profiles

Analysis of the flight profiles began with collecting the timestamps, latitudes, longitudes, and altitudes for the selected cruise leg. To reduce excess noise encountered when processing very close data points, every other point was used, resulting in spacing of approximately 2 minutes between each ETMS observation. These discrete steps in the flight path data are referred to here on as segments. From this information, the distance between each segment was calculated, along with the groundspeed components (north and east) based on the timestamps.

Groundspeed was then smoothed using a moving average over 7 data points (equating to approximately 15 min) to eliminate noise based on inaccuracies in the location and timestamp information.

Next, the nearest wind components and temperature data were collected for each point in the flight, based on time, altitude, and location. The process for searching the weather data files for particular points was itself complex. This involved first selecting the correct weather file based on time, decoding the location information from a Lambert conformal grid into readily understandable coordinates, and then utilizing binary search algorithms to find the point in the weather data nearest to the aircraft location. This process was repeated for each data point in the flight profile. Given the groundspeed and wind speed components, the airspeed profile was easily calculated as the vector difference between the two. Mach numbers were calculated using the airspeed magnitude and the local temperature.

Next, the vertical flight path angle was calculated for each segment. This was calculated knowing the air distance flown (airspeed multiplied by time) and the vertical altitude change. Based on the aircraft type, Mach, altitude, weight, temperature, and flight path angle, the SAR was calculated using the aforementioned method. Finally, the SAR was corrected for wind to produce the SGR metric, using the calculated ratio of groundspeed to airspeed. Each segment's fuel burn was calculated based on the local SGR and segment length. This fuel burn was subtracted from the aircraft weight to update the aircraft weight for the next segment. This procedure was completed in series for all segments in each flight path. A block diagram of this process is shown in Figure 11.



Figure 11. Diagram showing the process for fuel burn calculation of an existing flight profile.

The result of this fuel calculation procedure was a weight profile of the aircraft throughout the flight. The total fuel burn for the cruise leg was them simply the difference between the starting and ending weights. Each flight analyzed was calculated in this manner, yielding the baseline fuel burn to which improvements are compared.

3.3.5 Developing the Optimum Speed and Altitude Profile

Perhaps the most challenging step in the analysis was creation of optimum speed and altitude profiles. Each optimized trajectory was based on an actual one, leaving the original lateral path profile unchanged. The initial weight assumption for each of these modified trajectories is also considered identical to that of the base flight, which has been previously discussed. The only aspects changed were the speeds and altitudes along the path. In the initial iteration of the optimized trajectory, the same wind and temperature gathered for the original flight was used. Using the winds and temps at the optimized points was not yet possible because the ideal path had not been generated. In a second iteration, the winds and temperatures were found for the first optimized trajectory, and then subsequently used to generate a second iteration of the optimized trajectory which accounted for the correct wind and temperature.

The process started with selecting the optimum initial altitude at the initial assumed weight. Fuel burn was calculated for each segment and used to determine the altitude at the next

segment. This fuel burn calculation required an iterative solution because as fuel burn changes, the ideal altitude changes, which affects the flight path angle, which in turn has an effect on fuel burn. Additionally, the optimization of speed was iterative, because the optimal speed changes with wind and flight path angle. These interacting properties made the optimization challenging.

For each segment, after the weight and altitude were established from the previous segment, a speed optimization procedure was started. This process started with selecting the ideal Mach number from the SAR data, which essentially served as a first estimate of best Mach. After converting this Mach to airspeed based on the local temperature, the ground speed was calculated using vector algebra based on the wind components and aircraft airspeed. Next, an inner loop was created to converge on an optimal flight path angle, because SGR for each segment was circularly dependent on flight path angle. In this loop, SAR was estimated based on the current Mach estimate, aircraft weight, altitude, and flight path angle, then corrected for wind using the ratio of airspeed and groundspeed. The resulting SGR was used to calculate the fuel burn over the segment. This fuel burn was used to generate a new aircraft weight for the following segment, corresponding with a new ideal altitude and new estimate flight path angle. This process repeated until no change in fuel burn was observed and thus a solution converged. Returning to the Mach optimization loop, the converged fuel burn was stored with the first Mach estimate. The Mach loop then continued with a new speed estimate, simply perturbing the initial estimate both up and down. This process was repeated until the Mach yielding the minimum fuel burn was found. Based on this fuel burn, the aircraft weight was finally updated for the next segment, and a new ideal altitude calculated based on the weight versus altitude trend for the aircraft. The latter procedure was used for each segment of the flight, yielding altitude, speed, SAR, and weight profiles. A block diagram of this process is shown in Figure 12.



Figure 12. Diagram showing the process of generating the optimum speed and altitude profiles.

At the end of each cruise climb, a short descent was included to bring the aircraft down to a lower cruise altitude to negate any vertical potential energy differences between the original and optimized trajectory. The termination of the descent was chosen at an altitude such that the difference between the start and end altitudes of the optimized profile was identical to the difference between the start and end altitudes of the original profile. This descent segment played an important role in the calculation of the optimum altitude profile. During a cruise climb, the aircraft climbs for the entirety of the flight, burning extra fuel to climb and maintain the optimum altitude. A descent must be included to account for this energy that is ultimately recouped upon final descent anyway. An example of a cruise climb with descent compared with an original flat profile is shown in Figure 13. This figure demonstrates how the optimum and original profiles both start and end at the same altitudes and how both cover the same distance. Therefore, the performance comparison is totally fair.



Figure 13. A flat altitude profile (blue) and a corresponding cruise climb profile with descent (green).

The previous explanation of the optimal flight path calculation specifically explained the process of generating the completely optimal trajectory, optimizing both speed and altitude. These procedures were slightly modified to generate other partially optimal flight profile combinations: optimal speed, but unaltered altitude and optimal altitude, but unaltered speed. In the speed only optimizations case, the altitude profile was simply copied from the actual flight path, but the speed optimizer loop was still utilized to find an ideal speed at each point. For the altitude-only case, the speed profile was copied, but the altitude was updated at each segment based on the fuel burn.

3.3.6 Developing the Step Climb Profiles

The step climb cases were created by starting the flight at the initially assumed weight and altitude. The speed profile was simply copied from the original flight such that only altitude effects were visible in the results. As the aircraft traveled along level at the original altitude, fuel was burned and the ideal altitude was tracked accordingly. When the ideal altitude reached a height equal to half of the step distance, the aircraft began to climb at a nominal 0.5 degree climb angle. This is a representative climb angle based on observations of actual aircraft step climbs in the flight data. Once the altitude profile had climbed to the next step altitude, it again

continued in a level trajectory until due for anther climb. Like the other ideal profiles, a descent was included to match the potential energy change of the basis flight. An example output of the procedure for a 1,000 ft step climb is shown in Figure 14.



Figure 14. Example of 1,000 ft step climb trajectory generation, shown with the actual flight and optimal cruise climb trajectory.

Like the optimal altitude trajectory generation, a first iteration was generated using the same wind and temperature data found for the original trajectory. Because these values were invalid for the altitudes in this new step climb trajectory, a new set of wind and temperature data was then obtained for the initial step climb iteration. Then, a second and final trajectory iteration was calculated using these updated atmospheric conditions, such that the winds and temperatures matched the altitudes in the step climb. The rest of the process for calculating the fuel burn on these trajectories was identical to the rest of the analysis: the climbs and descents had negative and positive effects on the SAR, respectively, the SAR was adjusted for winds to create SGR, and the fuel burn rate was integrated to find the total fuel consumption for the step trajectory. This value was then compared with the fuel estimate for the original trajectory to gauge the potential benefit of the step climbs.

3.3.7 Developing the LRC Speed Profiles

The final profile case under consideration was a speed profile held constant at Long Range Cruise (LRC) speed, and an altitude profile unchanged from the original flight. The creation process of this profile was simply to find the LRC Mach, hold the speed constant at that number, and fly the original altitude profile. The LRC Mach was found in Piano-X by selecting the "Long Range Cruise" speed option, flying the aircraft on a detailed mission profile, and observing the selected Mach number at level cruise conditions. This process showed that the LRC cruise Mach did not vary across a range of altitudes, allowing a constant LRC Mach to be chosen with confidence. A profile was then created, forcing the speed to remain constant at that Mach number, and using the same altitude, wind, and temperature profiles found for the original flight. An example of the LRC speed profile, along with the original and fuel-optimal profiles, is shown in Figure 15.



Figure 15. Example of LRC Mach speed profile, along with the original and fuel-optimal profiles.

3.4 Regional Jet and Turboprop Analysis

The CRJ and the Q400 could not be included in the aforementioned flight analysis process for several reasons. After observing the flight data, it was clear that most of the flights were spent climbing and descending. Because of the short nature of these commuter flights, little or no time was spent at the cruise altitude, making a cruise analysis impossible or invalid. Any fuel savings calculated would only apply to a very small fraction of the total fuel burn, therefore having little effect on the trajectory performance. The short cruise also meant that the optimal altitude only increased by a few hundred feet over the entirety of the cruise, essentially rendering all step climb cases void. Additionally, the cruise altitudes were lower than even the lowest possible optimal altitudes for these aircraft types. This made weight estimation via the filed altitudes impossible, and any calculations of off-optimality for these fights required an alternate weight estimation method. Given these challenges, these aircraft were excluded from the detailed flight analysis.

These flights were examined more broadly by using a different weight assumption and simply observing the approximate difference between the operated altitudes and optimal altitudes based on the weight assumption. By definition, an aircraft can take off at no more than 100% MTOW. Assuming a conservative 5% reduction in weight due to fuel burned in takeoff and climb (a large portion of the CRJ and Q400 flights), it is highly unlikely that any of these aircraft were at less than 95% MTOW at the top of climb. In fact, most were probably at weights much less than that, especially considering few aircraft depart at maximum takeoff weight. According to the weight versus optimal altitude curve, a weight equal to 95% of MTOW corresponds to an optimal altitude of FL388 for the CRJ, the lowest optimal altitude conceivable, and yet still nearly 10,000 ft higher than the altitude filed (or reached) by most of the CRJ flights. Nearly all Q400 flights achieved an altitude of only 13,000-15,000 ft, when the optimal predicted altitude for that aircraft at 95% MTOW is FL340. To provide at least some look at the off-optimum altitude performance for these flights, the altitudes of the cruise legs (although short in duration) were compared with the minimum ideal altitudes of FL388 for the CRJ and FL340 for the Q400. The offsets between the actual cruise altitudes and these estimated optimum altitudes were referenced to the SAR data contours to obtain estimates of the potential fuel savings available if the optimal altitude had been reached.

A caveat of this analysis method is that because the flights are so short, any attempt to climb higher to the optimal altitude would lengthen the climb and descent phases, further shortening the cruise phase. In fact, reaching the optimal altitude would be impossible in many of these flights, as the climb to altitude would be interrupted by a descent. Therefore, these results are meant only to provide a general reference of the current cruise performance relative to optimum conditions, and do not necessarily imply attainable performance targets.

This broader examination was performed on a majority of the regional flights that were unfit for a detailed cruise analysis. However, a small subset of CRJ flights was found with cruise legs long enough to warrant processing via the detailed analysis tool. Twelve CRJ flights between Los Angeles and Salt Lake City were over 500 nm in length, and spent over 200 nm in cruise. The altitude and speed performance of these flights were calculated using the detailed flight analysis tool. These results were kept separate from the larger flight collection because the low cruise altitudes necessitated that the alternate 95% MTOW weight assumption be used.

Chapter 4

Results

4.1 Analysis Tool Output

Each flight was analyzed in detail using the process and MATLAB tool described in Chapter 3. A sample output from this tool is shown in Figure 16. Time, altitude, airspeed, groundspeed, Mach, SGR, weight, distance, latitudes, longitudes, wind components, and temperature were captured or calculated for each flight profile case, and multiple profile cases were analyzed for each flight. This depiction shows four of the major parameters: relative wind, expressed as the difference between ground and air speed, altitude, instantaneous SAR, and speed. The blue lines correspond to the unmodified flight as it was flown; the green lines correspond to the full optimal case (ideal speed and altitude). In this sample, notice that the ideal altitude profile descends and terminates at the same altitude as the original profile. This is required for a fair comparison between profiles to match the potential energy change during cruise as described in Section 3.3.5. The SAR profile shows that, as with most flights, the optimal SAR is steady and generally slightly higher than the observed SAR. The small SAR spikes are caused by noise in the groundspeed data or wind affecting the calculated airspeed, and the sharp SAR spikes indicate increases and decreases in fuel burn due to climbs and descents, respectively. As expected, the SAR on both profiles increases steadily over time as the weight of the aircraft decreases. Finally, the speed profile is typical of what was observed in the data, showing that the

flight operated at a speed higher than its fuel optimal value. Small perturbations in the speed results were due mostly to the granularity of the ETMS latitude and longitude points, which caused in fluctuations in the groundspeed calculations. The larger variations in speed were most likely due to the NOAA wind estimates that may or may not have accurately represented the actual winds encountered by the aircraft. The subtle changes in the optimal speed profile are the result of changing winds en-route, where higher relative headwinds resulted in a faster optimal speed, and vice versa.



Figure 16. Example of the many outputs of the performance analyzer tool, showing actual (blue) and optimal (green) profiles. This figure shows the winds of the flight, expressed as groundspeed minus airspeed (top left); altitude (top right); SAR (bottom left); and Mach number (bottom right).

While most flights contained cruise legs that began at the same altitude filed on the flight plan, some did not. Figure 17 illustrates such an example, where the filed altitude (FL300) was lower than the initial cruise altitude (FL320). In keeping consistency with the weight estimation assumption, the aircraft weight was assumed to be optimized at the initially filed altitude. The practical result of this is clearly visible—the actual flight profile appears to be above its optimal for the entirety of the flight, despite a step climb that indicates the FMS may have had weight data to the contrary. Therefore it is possible that in this example, the weight assumption resulted in the flight seemingly operating above optimal, thus causing incorrectly high fuel burn estimation. However, the weight estimate on most flights was such that the start of cruise was assumed to be at the optimal altitude, as shown in Figure 18. This is a conservative assumption, as it is unlikely that most flights are so close to their optimal altitude at start of cruise. As a result, many flights likely appeared to have less room for improvement than they actually did, and the altitude improvement potential results are therefore on the conservative side.



Figure 17. Actual (blue) and optimal (green) profiles for an MD82 from Washington DC to Dallas.



Figure 18. Actual (blue) and optimal (green) profiles for an A320 from New York to Los Angeles.

More than just optimal altitude and speed profiles were generated for each flight. As discussed in the methodology, optimal step altitude cases were analyzed to examine the potential benefit of step climbs, which are much more feasible to implement in the near term when compared with cruise climbs. An example of the various altitude optimization cases is shown in Figure 19, which shows a Boeing 757-200 flight from Los Angeles to New York. These profile types were generated for each flight, and the fuel burn of each compared to the original profile. Some of the shorter flights had step profiles that were level, because the aircraft never burned enough fuel to warrant a step up.



Figure 19. Example of the actual and improved altitude profiles for a Boeing 757-200 flight from Los Angeles to New York. Profiles like this were created for each flight.

In addition to the altitude cases, two improved speed profiles were generated: the optimal or best fuel economy case; and the constant Long Range Cruise (LRC) Mach case. As expected, the LRC speed was always faster than the fuel-optimal speed. In almost all flights, the operated speed was faster than best economy, and even faster than LRC in many cases. An example of these profiles is shown in Figure 20, for the same 757 flight.



Figure 20. Example of the actual and improved speed profiles for the same Boeing 757-200 from Los Angeles to New York. These speed profiles were created for each flight.

4.2 Individual Results

The primary objective of this entire exercise was to characterize the potential reduction in fuel burn that could be achieved on today's flights using only improved speed and altitude trajectories. A total of 257 selected flights were processed to calculate the total cruise fuel consumption before and after improvements were made, and then compared. The primary metric for comparison used was the percentage reduction in fuel burn achieved on a modified flight profile in comparison to the original.

This potential fuel burn reduction for each flight was captured graphically on several plots, organized by trajectory type, flight distance, aircraft type, and airline. This graphic illustration attempts to display all results in an easily understandable manner.

Figure 21 shows the maximum potential fuel burn reduction for each flight analyzed. The yaxis represents the total possible fuel saving for that flight, using ideal altitude and speed profiles. The x-axis shows the cumulative distance flown in the cruise leg for each flight. By showing the cruise distance instead of the stage length or total cumulative distance, this representation clarifies how much of distance of each flight was actually being used for the fuel burn analysis. The horizontal line represents the aggregate fuel burn savings over all flights. The clustering of points is due to the selection of the city pairs, which dictated certain common distances.



Figure 21. Maximum possible fuel savings for each and every flight analyzed, categorized by aircraft type.

To further clarify the results, this information was split up into different plots by aircraft type, with flights in each plot categorized by the airlines which operated the flights. Figure 22 through Figure 25 show plots of fuel savings versus cruise distance separated by aircraft type for only the optimal speed and optimal altitude results. In these plots, each flight is represented by a single point. The horizontal lines represent the fuel-weighted averages for each aircraft type, which are later listed numerically in Table 4.



Figure 22. Maximum possible fuel saving potential for Airbus A320-200 flights. The plot on the right shows different shapes and colors for the different airlines.



Figure 23. Maximum possible fuel saving potential for Boeing 737-700 flights. The plot on the right shows different shapes and colors for the different airlines.



Figure 24. Maximum possible fuel saving potential for Boeing 757-200 flights. The plot on the right shows different shapes and colors for the different airlines.



Figure 25. Maximum possible fuel saving potential for McDonnell Douglas MD82 flights. The plot on the right shows different shapes and colors for the different airlines.

The maximum possible benefit plots are useful in placing bounds on the calculations, but provide little insight into which forces are driving those results. Figure 26 shows the fuel saving potential if either speed or altitude was optimized, but the other was not. For each flight, the yaxis value represents the fuel saving that resulted from optimizing the speed but leaving altitude untouched, and the x-axis represents fuel savings resulting from optimizing altitude but not speed. Points on the high side of the diagonal represent flights that benefited more from speed changes than altitude changes, and vice versa for the low side. The percentages in the top right corner depict the fraction of points above or below this line. Some points have negative altitude optimal (x-axis) points. In these cases, the fuel burn in the cruise climb case was actually worse than the unaltered case, which occurred when increased headwinds on the improved altitude profile outweighed the performance benefit of the cruise climb. Because results were calculated for flights going both directions on the chosen routes, this wind also conversely provided positive benefit in other cases, mostly averaging out the wind effect. Figure 27 through Figure 30 show the same data, but separated by aircraft type.



Figure 26. Optimum speed-only versus optimum altitude-only fuel burn reduction potential for all flights, categorized by aircraft type.



Figure 27. Optimum speed-only versus optimum altitude-only fuel burn reduction potential for Airbus A320-200, grouped by airline.



Figure 28. Optimum speed-only versus optimum altitude-only fuel burn reduction potential for Boeing 737-700, grouped by airline.



Figure 29. Optimum speed-only versus optimum altitude-only fuel burn reduction potential for Boeing 757-200, grouped by airline.



Figure 30. Optimum speed-only versus optimum altitude-only fuel burn reduction potential for McDonnell Douglas MD82, grouped by airline.

Optimal cruise climbs represent the best case scenario for altitude profiles, but step climbs are much more practical given the system limitations. RVSM permits vertical separation as low as 1,000 ft, requiring steps of 2,000 ft on two-way airways and 1,000 ft on one-way airways. The step climbs follow the ideal altitude as best as possible, starting at the filed altitudes, which are equivalent to the starting altitudes of the full optimal cases as well. Figure 31 and Figure 32 show the cruise fuel burn reduction if speed was left unchanged from the original flight, but altitude was modified to 1,000 ft step climbs. They first is categorized by aircraft type, and the second, by airline. Figure 33 and Figure 34 show the same results but for 2,000 ft step climbs. The negative values again indicate cases where the step climb trajectory encountered headwinds high enough to negate the altitude benefit and actually reduce the fuel burn. The gray line represents the fuel-weighted average across all flights.



Figure 31. Savings achieved by implementing 1,000 ft step climbs and leaving speed unchanged, categorized by aircraft type.



Figure 32. Savings achieved by implementing 1,000 ft step climbs and leaving speed unchanged, categorized by airline.



Figure 33. Savings achieved by implementing 2,000 ft step climbs and leaving speed unchanged, categorized by type.



Figure 34. Savings achieved by implementing 2,000 ft step climbs and leaving speed unchanged, categorized by airline.

The optimal speed results show significant benefit. However this optimal speed is a best economy setting, which is inevitably slower than most carriers would choose to fly even if any speed were available. The long range cruise (LRC) setting is a faster speed yet still fairly efficient, representing a likely selection by operators when time is critical to cost. Although it corresponds to a high CI, it still offers an efficiency improvement over many of the flights analyzed which exhibited very high speeds. Figure 35 and Figure 36 show the potential fuel burn reduction in cruise for each flight if speed was adjusted to be constant at the LRC setting, and the altitude profile was left unaltered. The first groups the results by aircraft type, and the second, by airline.



Figure 35. Savings achieved by adjusting speed to the LRC setting and leaving altitude unchanged, categorized by aircraft type.



Figure 36. Savings achieved by adjusting speed to the LRC setting and leaving altitude unchanged, categorized by airline.

4.3 Aggregate Results

The results were first aggregated by aircraft type to highlight any significant differences in the operational improvement potential between types. For an identical set of flight profiles, this difference corresponds to the sensitivity of a given aircraft type to off-optimal conditions. However, because each flight analyzed was unique, these results may include effects beyond the performance of the aircraft, including varying airline operating standards and specific route characteristics. By including these real effects, these results provide the most realistic improvement potential estimates possible.

Table 4 shows the aggregate fuel burn saving potential for each aircraft flying certain enhanced altitude and speed profiles. To calculate these aggregates, the calculated fuel burn for each *original* flight of that aircraft type was summed together. Then the fuel burn of each corresponding *improved* trajectory was summed. The savings were then found by calculating percent change from the *original* fuel sum and the *improved* fuel sum. The practical implication of this method is that instead of the results being weighted equally per flight, the influence of each flight's result on the aggregate is proportional to the fuel burn of the flight. More simply, this result answers the question: If the trajectories for all flights of Airplane-X were improved to Profile-Y, what would be the net savings in fuel burn across all those flights?

Speed Profile	Optimal	Optimal	Unaltered	Unaltered	Unaltered	LRC
Altitude Profile	Optimal	Unaltered	Optimal	Step 1k ft	Step 2k ft	Unaltered
A320	3.43%	2.38%	1.75%	1.15%	0.96%	1.54%
B737	1.98%	1.17%	1.03%	0.98%	0.90%	-0.82%
B752	4.05%	2.69%	1.70%	1.42%	1.28%	2.13%
MD82	2.68%	2.26%	0.70%	0.20%	0.40%	1.59%
All Flights	3.48%	2.40%	1.50%	1.12%	1.02%	1.61%

Table 4. Cruise fuel savings potential by aircraft type.

The fuel burn reduction potential was also grouped by city pair to help identify any routespecific behavior differences that might be caused by congestion, airspace, or ATC issues unique to the given flight route. Table 5 shows the results for various profile improvement strategies, grouped by the route of flight. Details corresponding to these routes, including cruise leg length, total flight distance, and the cruise fraction of flight are shown in Table 6. This table is important because it shows how little of the flight can be usable on short routes. As the length of cruise becomes very small, the fuel burn in cruise represents a much smaller fraction of the total fuel burn, so the potential fuel saving calculations become less significant.

Speed Profile	Optimal	Optimal	Unaltered	Unaltered	Unaltered	LRC
Altitude Profile	Optimal	Unaltered	Optimal	Step 1k ft	Step 2k ft	Unaltered
Atlanta - Miami	4.76%	3.69%	2.10%	2.04%	2.05%	2.86%
Washington DC - Chicago	5.55%	3.57%	3.73%	3.11%	3.12%	2.52%
New York - Chicago	2.60%	1.94%	0.87%	0.43%	0.66%	0.97%
Washington DC - Dallas	2.99%	2.49%	0.86%	0.35%	0.45%	1.82%
Los Angeles - Chicago	3.62%	2.14%	1.51%	1.11%	0.99%	1.18%
Los Angeles - New York	3.59%	2.50%	1.73%	1.35%	1.16%	1.76%
Boston - San Francisco	3.31%	2.27%	1.07%	0.82%	0.61%	1.80%

Table 5. Cruise fuel savings potential by flight pair, sorted by distance.

 Table 6. Aggregate flight characteristics by city pair, corresponding to the results in Table 5.

City Pairs	Average Cruise Dist. (nm)	Average Total Dist. (nm)	Cruise Dist. Fraction	Flight Count
Atlanta - Miami	260	578	44.9%	27
Washington DC - Chicago	258	549	47.1%	19
New York - Chicago	375	678	55.3%	77
Washington DC - Dallas	776	1063	73.0%	25
Los Angeles - Chicago	1255	1590	78.9%	41
Los Angeles - New York	1919	2247	85.4%	60
Boston - San Francisco	2144	2448	87.6%	8

While ATC procedures and airspace limitations play a large part in limiting the success of achieving optimal cruise operations, a significant amount of responsibility is held by the airlines. In an attempt to help distinguish any differences in the efficiency of various airlines' procedures, the results were grouped by airline. The airline names were kept anonymous given the sensitivity of this information. Table 7 shows the fuel burn reduction potential of three profile optimization cases, grouped by airline. Table 8 shows the corresponding aggregate flight data statistics to help put the results in context.

Speed Profile	Optimal	Optimal	Unaltered	Unaltered	Unaltered	LRC
Altitude Profile	Optimal	Unaltered	Optimal	Step 1k ft	Step 2k ft	Unaltered
Airline 1	3.64%	3.52%	1.21%	1.51%	1.51%	1.13%
Airline 2	3.02%	2.37%	0.85%	0.42%	0.53%	1.75%
Airline 3	2.97%	2.21%	1.34%	1.11%	0.94%	0.96%
Airline 4	3.53%	2.23%	1.83%	1.63%	1.50%	1.66%
Airline 5	3.25%	2.32%	1.70%	1.08%	0.89%	1.52%
Airline 6	1.82%	1.13%	0.87%	0.83%	0.77%	-0.87%
Airline 7	4.26%	2.76%	1.75%	1.34%	1.22%	2.11%
Airline 8	3.37%	2.45%	1.91%	1.31%	1.01%	1.64%

Table 7. Cruise fuel savings potential by airline.

 Table 8. Aggregate flight characteristics by airline, corresponding to the results in Table 7.

Airline	Average Cruise Dist. (nm)	Average Total Dist. (nm)	Cruise Dist. Fraction	Flight Count
Airline 1	252	572	44.0%	5
Airline 2	673	968	69.5%	65
Airline 3	1482	1768	83.8%	14
Airline 4	909	1225	74.2%	36
Airline 5	1516	1864	81.3%	23
Airline 6	835	1145	72.9%	22
Airline 7	891	1209	73.7%	82
Airline 8	1907	2252	84.7%	10

4.4 Regional Jet and Turboprop

The regional jet (CRJ-200) and turboprop (Dash 8 Q400) flights were examined separately due the brevity or absence of the cruise phases, and because low operating altitudes necessitated an alternate weight estimation process.

4.4.1 Regional Jet Results

Most of the regional jet flights were no more than 300 nm in length, and with cruise legs of only 50-100 nm. Figure 37 through Figure 39 show three representative CRJ flights: Los Angeles to San Francisco, Pittsburgh to New York, and New York to Washington DC. The blue lines depict the altitude of the aircraft, and the green lines represent filed altitude, which sometimes changes en route as flight plans are amended. These depictions demonstrate the brevity of the cruise legs.



Figure 37. CRJ-200 flight from Los Angeles to San Francisco. The green line represents filed altitudes in the flight plan.



Figure 38. CRJ-200 flight from Pittsburgh to New York. The green line represents filed altitudes in the flight plan.



Figure 39. CRJ-200 flight from New York to Washington DC. The green line represents filed altitudes in the flight plan.

The altitudes of the cruise legs for these flights were recorded and compared with the ideal altitudes generated by the 95% MTOW weight assumption. This altitude delta was then applied to the aircraft performance SAR contours to provide insight into the typically off-optimality of the CRJ-200 in cruise. The results for several city pairs are shown in Table 9.

 Table 9. CRJ-200 performance penalty estimate due to low cruise altitudes, relative to the optimal altitude of FL388.

City Pair	Altitude Reached	Fuel burn performance penalty due to off-optimum altitude	
Los Angeles - San Francisco	FL290-310	10-14%	
New York - Washington DC	FL170-210	30-37%	
New York - Pittsburgh	FL250-270	18-23%	

Despite the apparent futility of improving the cruise trajectories of these short regional flights, a small subset of the CRJ flights was found with significantly longer cruise legs. An example of these flights is shown in Figure 40.



Figure 40. CRJ-200 flight from Los Angeles to Salt Lake City. The green line represents filed altitudes in the flight plan.

These CRJ flight had cruise legs of 40-60% of the entire flight, enough to warrant processing via the flight analysis tool. Due to the small subset of these flights, and the alternate weight estimation method, they were considered separately from the other 257 flights. A total of 12 CRJ-200 flights between LA and Salt Lake City were analyzed, using the 95% MTOW assumption used previously. Examples of the altitude and speed profiles generated by the analysis are shown in Figure 41 and Figure 42. The aggregate numerical results are shown in Table 10.



Figure 41. Altitude profiles generated for CRJ-200 flight from Los Angeles to Salt Lake City.



Figure 42. Speed profiles generated for CRJ-200 flight from Los Angeles to Salt Lake City.

Table 10. Cruise fuel savings potential results of 12 CRJ-200 flights between Los Angeles and SaltLake City.

Speed Profile	Optimal	Optimal	Unaltered	Unaltered	Unaltered	LRC
Altitude Profile	Optimal	Unaltered	Optimal	Step 1k ft	Step 2k ft	Unaltered
CRJ-200, LA to SLC	6.55%	8.76%	5.65%	7.48%	7.46%	7.23%
CRJ-200, SLC to LA	3.14%	1.17%	5.91%	0.54%	0.58%	0.70%
CRJ-200, LA-SLC both ways	5.67%	6.79%	5.72%	5.68%	5.68%	5.54%
4.4.2 Turboprop Results

The Dash 8 Q400 turboprop flights had characteristics similar to those of the regional jet. These flights were short – less than 300 nm – and also had relatively short cruise legs, mostly 50 nm or less. Two examples are shown in Figure 43and Figure 44 for flights from New York to Washington DC, and Pittsburgh to New York, respectively.



Figure 43. Q400 turboprop flight from New York to Washington DC. The green line represents filed altitudes in the flight plan.



Figure 44. Q400 turboprop flight from Pittsburgh to New York. The green line represents filed altitudes in the flight plan.

Like many of the CRJ flights, the cruise legs for the Q400 were considered too short for a detailed analysis to be meaningful, and the filed altitudes are roughly half that of a nominal optimal altitude for this aircraft. Using the same 95% MTOW assumption as the maximum possible weight at top of climb, the ideal altitude for this weight is FL340. This altitude was used as a reference point to gauge how far from optimal the aircraft were operating. An estimate of the performance penalty caused by being far below the optimal altitude is shown in Table 11.

 Table 11. Q400 performance penalty estimate due to low cruise altitudes, relative to the optimum altitude of FL340.

City Pair	Altitude Reached	Fuel burn performance penalty due to off- optimum altitude	
New York - Washington DC	FL130-160	31-36%	
Pittsburgh - New York	FL160-170	28-31%	

Much like the CRJ flights, the caveat to these results is that given the short duration of the flight, reaching the optimal altitude would be impossible. Therefore these results stand as a general reference to the extent at which many turboprops operate below optimum altitude.

4.5 System-wide Benefit Potential

The cruise fuel burn savings results for the selected set of flights were used to develop an estimate of the total system-wide benefit potential. According to the Bureau of Transportation Statistics, 11.34 billion gallons of jet fuel were consumed in 2009 by certificated air carriers operating domestically (BTS, 2010). Assuming all was burned, 122.5 million tons of CO_2 were emitted by those operations. These values were used as starting points to estimate the absolute amount of fuel and CO_2 that could be saved in the US if certain speed and altitude strategies were implemented.

The results of this analysis provide fuel savings potential of the cruise leg of flight only. Therefore, corrections were made to estimate the overall fuel savings. The distance flown in cruise and the total distance covered for each flight were known for each flight from the path data. A basic approximation was made that fuel fraction of the flight equaled the distance fraction. That is, if the cruse phase occurred over half the total flight distance, then the cruise phase burned approximately half of the total fuel. The reasoning is that the cruise phase burns more fuel per distance than cruise, but the descent burns proportionally less as potential energy is recovered. This is a simple assumption that holds true as long as the engine and aerodynamic performance in climbs or descents does not vary dramatically from that in cruise. For the purpose of this calculation, this assumption was appropriate.

The results of the analysis were then grouped by flight distance, and the corresponding average cruise distance fraction for each group calculated. These values are listed in the third column of Table 12. As expected, the longer flights spent more time in cruise and had a larger cruise distance fraction; the opposite was true for short flights. Therefore, the cruise savings potential results for the longer flights had a larger effect on the total fuel burn than the shorter flights. The estimated cruise fuel consumption for all flights were aggregated by distance group, creating a total group cruise fuel burn (fourth column). This number was then adjusted by the cruise distance fraction to provide an estimate of the total flight fuel burn of that group (fifth column) utilizing the flight phase fuel contribution assumption above. These totals for the groups were then normalized by the total fuel burn of all flights to develop relative fractions of each group's contribution to the total fuel burn (sixth column).

 Table 12. Cruise distance fractions and calculated weightings for the relative contribution of each

 distance group to the total fuel burn of all flights.

Flight Distance Traveled (nm)	Number of Flights	Average Cruise Distance Fraction	Total Cruise Fuel Burn (thousand kg)	Estimated Total Flight Fuel Burn (thousand kg)	Estimated Total Fuel Burn Fraction
Less than 1000	123	51.7%	254	493	22.9%
Between 1000 and 1500	25	73.0%	140	192	8.9%
Between 1500 and 2000	41	78.9%	324	410	19.1%
Greater than 2000	68	85.7%	905	1056	49.1%
All flights	257	75.18%	1623	2151	100%

The fuel savings results for each operational strategy were corrected via the cruise distance fraction to develop savings percentages relative to the entire flight instead of just cruise. These flight fuel savings results were then weighted by the total fuel consumption of each group (sixth column of Table 12) to develop the overall fuel reduction potential of each operational strategy. These were then directly applicable to the system-wide total fuel consumption and CO_2 emission values above. The calculated system-wide benefit of each strategy is shown in Table 13.

Speed Profile	Optimal	Optimal	Unaltered	Unaltered	Unaltered	LRC
Altitude Profile	Optimal	Unaltered	Optimal	Step 1k ft	Step 2k ft	Unaltered
Cruise Benefit Potential	3.48%	2.40%	1.50%	1.12%	1.02%	1.61%
Total Benefit Potential	2.63%	1.82%	1.15%	0.86%	0.79%	1.22%
Fuel Reduction Potential (million gallons per year)	298.4	206.0	130.2	97.4	89.2	138.2
<i>CO</i> ₂ Reduction Potential (million tons per year)	3.22	2.23	1.41	1.05	0.96	1.49

 Table 13. Total system-wide fuel burn and CO2 reduction potential.

Chapter 5

Discussion of Results

5.1 Maximum Improvement Potential

The results of this analysis put an upper bound on the benefits that can be achieved via altitude and speed optimization. One of the most important pieces of information gathered was the magnitude of fuel burn reduction possible using simply operational changes. The average potential available to reduce fuel consumption in cruise across all flights was 3.5%, or 2.6% when referenced to the total flight fuel burn. This is an impressive result considering that nearly all flights can benefit from operational improvements. Additionally, cruise phase consumes the majority of the fuel for most flights, which makes these results even more meaningful in the context of system-wide potential. A 3% reduction in cruise fuel burn would be an impressive target even for major aircraft upgrades like engine or aerodynamics enhancements, and yet no new aircraft are needed to implement improved altitude and speed profiles.

Another major result of this analysis is the variability in improvement potential seen across the flights processed. Flights demonstrated savings ranging from nearly none up to 12% in some cases. This fact alone shows that how an aircraft is operated can have a dramatic affect on its fuel burn. Observations of the flight data showed that some flights were operated much more efficiently than others, so the results were truly unique for each flight. The sensitivity to operations highlights the large responsibility carried by those who dictate flight operations, including pilots, airlines, and ATC.

The variability in improvement seemed to be heightened at smaller ranges. On flights where the cruise leg was less than 500 nm, improvement percentages exhibited a much wider spread in magnitude, and a higher averages overall. One reason for this is likely due to the neglecting of performance goals on short flights. On shorter flights, pilots have less distance over which to make up time for delays; this problem likely contributed to these shorter flights being flown at fast inefficient speeds. Additionally, altitude optimization is likely deemed less critical on short flights due to the lower absolute quantity of fuel burned, so operators may be less likely to seek ideal altitudes. Likewise, while some of these potential savings appear high for short flights, the absolute benefit is much less significant than for longer ones.

Examining the results by aircraft type, one sees that the B752 represented the highest potential for overall improvement and the B737 the lowest. Part of this difference is probably rooted in the various distances and routes flown by the two aircraft, as the B752 was operated over longer ranges on average. However one of the more influential contributors is the type of airline which operates each aircraft. The B752 is an older aircraft, flown mostly by older legacy carriers which put less emphasis on the details of cruise efficiency. The 737-700, however, is the newer "next generation" variant of the popular aircraft, and is favored by newer carriers whose priority is reducing cost.

5.2 Altitude Effects

5.2.1 Optimum Altitudes

The aggregate cruise fuel savings potential for complete altitude optimization, with no speed changes, is 1.5%. With the exception of the short routes (500-600 nm) and the few Boston to San Francisco flights, the average altitude benefit from ideal profiles increased with range. This

behavior is likely due to the fact that on longer range flights, more fuel is burned and the ideal altitude changes more dramatically. As evidenced by the flight data, step-ups in altitude were sometimes made but often profiles were flat at best, even on the longer New York to Los Angeles route. These flat profiles caused increased detriment to the performance as the flights progressed. As expected, this result shows that altitude optimization is especially important on longer flights, where the ideal altitude changes more dramatically and a larger amount of fuel is being burned. The unusually high fuel reduction potential for the 500-600 nm routes is likely due to the previously mentioned idea that fuel efficiency is an afterthought on short flights where the return on such operational procedures are low in comparison to the entire fuel burn of the flight.

The altitude-only optimization results varied by airline, averaging savings 1-2%. This variation was likely in part due to the differing route lengths flown by the airlines, but also may have reflected unique company policy and training procedures. Indeed, one of the newer low-cost carriers showed approximately half as much altitude savings potential compared to some of the older legacy carriers. This trend would imply that lower cost carriers are taking more care in the flight planning and operation stages to maximize flight efficiency.

An examination of the aggregate results by type shows that the A320 and B752 command the highest benefit potential from altitude optimization. A closer look at the SAR sensitivity contours shows that the A320 does exhibit a significant drop-off in performance at altitudes above optimal and at higher speeds, which may contribute to the altitude-based fuel burn penalty it sees in today's operations. The B752, however, also currently exhibits performance loss due to off-optimum altitude conditions but its SAR contour shows a relatively shallow slope. The reason for this discrepancy is likely due to the types of operations it sees. The B752 is larger than the A320 and more of its flights are flown at longer ranges, so it is apt to experience larger

differences in its operated and ideal altitude on long flights if the operators are not diligent in preventing this deviation.

The short regional jet and turboprop flights numerically show a large potential for improvement, ranging from 10-30%. However, these larger values only apply to short flights with correspondingly short cruise legs, where the ideal altitude could never be reached. Therefore the potential for improvement in cruise is little to none. Only improvements in the climb and descents can yield any real benefit on these flights; that analysis is out of the scope of this research. The longer range regional jet flights reached altitudes much closer to the optimal altitude, and showed more moderate potential fuel savings of 5.7% for the optimal altitude case. Given the size of these aircraft and short range in general, their potential contribution to fuel burn and emissions reduction overall is relatively small. Additionally, were these regional jets to climb higher to their optimal altitude, the climb and descent phases would have consumed even more of the short cruise leg, making the cruise benefit even less substantial.

5.2.2 Step Climbs

These step climb results generally scaled with the cruise climb results; that is, if one airline or type showed the most benefit potential from optimized altitude, it also showed the most potential from step climbs. The MD82 showed the lowest potential for altitude improvements, and unlike the general trend it showed less benefit from 1,000 ft step climbs than from 2,000 ft step climbs. The reason for this is again related to winds, where high altitude headwinds caused poorer performance on flights at higher altitudes. Many of these flights were short enough such that the aircraft never reached the point of stepping 2,000 ft, but still were able to step 1,000 ft. Therefore, the 2,000 ft profiles were lower and subjected to fewer headwinds. This explains the same unexpected trend in some of the results grouped by city pair. While considering flights in both directions sometimes minimized this wind effect, atmospheric conditions often change throughout the day, allowing return flights to face the same headwind as the outbound flights.

The total step climb benefit potential in cruise across all flights was 1.12% for 1,000 ft steps and 1.02% for 2,000 ft steps, compared with 1.5% for cruise climbs. As expected, the cruise climbs yield the most potential, and the largest step climbs yield the least potential. However, the data suggests that the benefit of the step climbs is still significant, and differs little from the more complex cruise climbs. The reason for this is simple. The altitude sensitivity to SAR for nearly all aircraft is fairly low in the vicinity of the optimal altitude. That is, deviations of 1,000 feet from the optimum often only incur a penalty of only 1% or less. Because both step climb profiles are generally within 1,000 ft of the optimum altitude at any point in the profile, their difference is fairly minimal. However, the absolute benefit of these profiles is still significant because the actual aircraft profiles are often far from any of the improved profiles. The regional jet results exaggerate this effect; in those, the benefit from step climbs (5.7% and 5.5% for 1,000 and 2,000 ft step climbs, respectively) is nearly identical to the benefit of a cruise climb. Here, the steps were often never made due to the short duration of the flights, and the cruise climb barely gained any altitude. All improved profiles were thus nearly identical, but yet far above the operated cruise altitude, thus resulting in the high magnitude of the savings potential. This and nearly all the other results demonstrate that any improvement in the altitude profiles, including simply starting near the optimal, is likely to result in fuel savings.

These results make a compelling case for step climbs, which can be implemented today in 2,000 ft steps under RVSM, traffic and weather permitting. Still, the choice of when or if to step climb may not always be obvious to pilot, considering wind conditions at the next altitude could be unfavorable, and will remain unknown until the aircraft has reached that altitude. Better flight planning and utilization of today's step climb capability is likely to yield considerable fuel savings when compared to current operations.

5.3 Speed Effects

5.3.1 Optimal Speed

Probably the most obvious trend visible in these results is that flights are generally much more susceptible to encountering off-optimal speed conditions than off-optimal conditions. In most flights, this is manifested in flying faster than the best fuel economy speed. In fact, 70% of all flights experienced more benefit potential from speed improvements than from altitude improvements. The aggregate speed-only potential savings range from 2-3%, consistently higher than the 1-2% figure seen with the altitude-only optimization. The SAR plots do indeed make clear that even speed changes by only a few hundredths of a Mach can reduce performance by several percent. These revelations are important because unlike altitude optimization, speed changes are more readily controlled by the pilot and less subject to restriction by controllers. These factors make cruise speed an easy target for improvement.

Results breakdown by aircraft type indicates that the B752 and MD82 show the most potential for improvement, while the B737 shows the least. The SAR contour plots make distinguishing differences due to speed sensitivity difficult at best. In fact, one might argue that the newer B737's performance at the cruise condition has been refined the furthest and thus would likely face higher penalties when operating off-optimum. However, the types of operators flying these aircraft again stand as a likely explanation. Only legacy carriers fly the MD82 and B752, while newer more efficient airlines feature the newer B737. The variation in their operating practices probably accounts for some of the difference in results.

Finally, the distribution by city pairs depicts a fairly consistent savings potential across various routes, except for the short routes. All but the shortest pairs exhibit speed-only fuel burn savings of 2-2.5%, versus 3.5% for the short flights. As mentioned previously, this may partly be explained by the brief nature of flights, which dictates that speed changes must be more dramatic if the time spent en-route is going to noticeably change. Therefore, any pressure

from dispatcher to meet on-time targets would be met with more significant changes in speed for shorter flights. Additionally, relative changes in fuel burn performance are less visible to pilots on shorter flights because the fuel-over-destination estimation by the FMS will not vary significantly with even moderate changes in performance efficiency. Alternatively, minor speed changes on long range flights may have drastic effects on the absolute fuel reserve calculations. Therefore, operators of short range flights may be less prone to observe the fuel efficiency sensitivity and may select faster cruise speeds.

A surprising result from the optimal speed profiles showed that optimal speed decreased in descents. This is a result of the correction from SAR in level flight to SAR in a climb or descent. In a descent, thrust is reduced, therefore reducing the total fuel burn. Although the equivalent level SAR for that speed is reduced, the descents adds fuel-free thrust, which increases the corrected SAR. Put more simply, one may envision the limiting case of a steep descent requiring zero engine thrust (engine off). In this case, the optimal speed is the best glide speed, which is lower than the max range cruise speed for jet aircraft. Hence, optimal speed is reduced as thrust is reduced in a descent.

5.3.2 LRC Speed

The benefits of constant Long Range Cruise (LRC) speed trajectories were less than that of fully optimized speed trajectories, which was expected as LRC is not fuel-optimal. LRC benefits scaled with optimal speed benefits in general, meaning that aircraft or city pairs exhibiting the most optimal speed benefit also showed the most LRC benefit. The reasoning for this effect is that the flights that showed the most potential for speed improvement generally flew at speeds much higher than optimal, and higher than LRC as well. Therefore the offset of the actual cruise speed from LRC was highest for these flights as well.

One aircraft type and one airline actually displayed an overall negative potential (increased fuel burn) for the LRC trajectories, but yet still had a positive benefit from complete speed

optimization. This is due to the fact that many of these flights operated at speeds below LRC such that while total speed optimization provided improvement, the LRC speed was higher and caused an increase in fuel burn. Not surprisingly, the aircraft with this result was the Boeing 737, and the airline was a newer low-cost carrier, indicating that this operator was flying at a more optimal speed below LRC. The aircraft and airlines with the highest benefit potential from flying at LRC were the Boeing 757 and MD82, flown by legacy carriers.

The aggregate cruise benefit across all flights for flying constant LRC was 1.6%, compared with 2.4% for total speed optimization. Despite this value being expectedly lower, it is still a significant amount when compared with other savings figures. In fact, this value is larger than all of the altitude-only improvement results. Therefore, if all aircraft flew a constant LRC but made no changes to their altitude profiles, the fuel savings would be greater than if aircraft flew cruise climbs and left speed unchanged. The conclusion is that speed has a larger effect on efficiency than altitude given the way that aircraft are operated today.

Chapter 6

Operational Barriers and Mitigations

6.1 Altitude Improvements

Operations like cruise climbs and step climbs can clearly reduce fuel burn and emissions, but they are plagued with system limitations and other barriers to their implementation. Some of these barriers and potential mitigation options are discussed below.

6.1.1 Cruise Climbs

Cruise climbs, if executed correctly, undoubtedly provide the best possible performance among altitude profiles (disregarding changing wind conditions with altitude). While the implementation of these operations is plagued with challenges, there are potential mitigations that can help reduce these barriers.

There are several primary challenges that need to be overcome before anything resembling a cruise climb can be made a reality. First and foremost, ATC is not equipped for or capable of safely permitting cruise climbs. Only lateral information is visually depicted on controllers' radar screens; vertical position and speed information is provided textually. This setup makes vertical separation a challenge, and thus only level flight is allowed for complexity reduction. Collisions are much easier to see and predict when only two lateral dimensions are in play; such predictions would be extremely difficult at best in three dimensions. To compound the problem,

inaccuracies with radar data cause controllers to apply large buffers on aircraft separation, which restrict aircraft movement. Airspace also poses a problem. Allowable altitudes are discretized by the thousands of feet. Most current airways are setup such that the direction of flight along the airway alternates with altitude. This allows controllers to handle bidirectional traffic and still have confidence in maintaining separation. With this arrangement, a continuously climbing aircraft would impinge on opposite-direction traffic flow, likely causing conflicts.

Future NextGen technologies do promise benefits that will help circumvent some of these issues. As more aircraft are equipped for higher levels of Required Navigation Performance (RNP), they will be able to maintain much tighter path tracking. This position information shared over Automatic Dependent Surveillance Broadcast (ADS-B) and make the global traffic picture more accurate, allowing ATC to potentially reduce separation minimums. Improved situational awareness will allow greater flexibility with aircraft movement, because controller uncertainty is reduced. A practical application is a multi-lane airway, where aircraft are placed at the same altitude and then spaced laterally based on tighter RNP minima. This airway would allow increased traffic capacity, which could expand the space available for aircraft to climb. In other words, one or more lanes could be devoted to climbing aircraft, which would help to reduce the potential of conflicts caused by the cruise climb. Also, ADS-B promises a shared common traffic picture between all flying aircraft, which could potentially be used to place some separation responsibility on the aircrafts themselves instead of being handled entirely by controlling facilities. This could negate the need for a more complex three dimensional situational awareness picture for the controllers, and simultaneously reduce controller workload.

Limitations with aircraft FMSs and other avionics systems pose a problem. Currently, avionics are not setup to handle cruise climbs. The FMS does track aircraft weight and can

suggest ideal flight levels, but cannot provide continuous high-resolution optimal altitude data. Additionally, autopilot systems provide a means of altitude hold in cruise but are not designed to handle continuous cruise climbs. Of course, the required information to perform these climbs is available, but avionics must be redesigned to take advantage of that information. New systems take money and time, which is an obvious obstacle. However, if the larger ATC barriers to cruise climbs were removed, and airlines showed interest in this capability, avionics manufacturers would surely follow.

6.1.2 Step Climbs

While not as effective at improving performance as cruise climbs, step climbs do provide an intermediate level solution to increasing altitude performance. Better yet, today's system largely supports step climbs in cruise flight. Traffic and workload permitting, controllers will grant altitude change requests. Still, analysis of the flight paths showed that opportunities to make step climbs were seemingly not always acted upon.

One factor contributing to the lack of observed step climbs may be lack of motivation on behalf of the pilot. Controllers have no performance information about the aircraft and are not responsible for aircrafts' operational efficiency, so pilots need to explicitly request altitude steps as needed. If the airline does not have a firm policy on maximizing the use of step climbs, many pilots are apt to ignore them, especially on shorter flights. The choice to make these operational efficiency improvements comes from the airline and not the FAA, so airlines may need to be more diligent with seeking optimal altitudes. In other cases, turbulence at altitude dictates a change in altitude from optimum. Passenger comfort always comes at a higher priority than step climb efficiency, and issues like this will likely remain unsolved in the near term.

Still, some of the greatest barriers to optimal step climb profiles are ATC and traffic congestion. ATC will sometimes force traffic to climb or descent to avoid conflicts; these deviations disturb the altitude performance. With technologies like ADS-B and RNP, however,

airway capacity could be increased, reducing the need for altitude deviations. The improved accuracy could also allow for more efficient passing options which would reduce deviations caused by overtaking aircraft. RVSM has helped to increase step climb efficiency by reducing steps from 4,000 ft to 2,000 ft in size, allowing aircraft to be much closer to their ideal altitude. If RVSM can be coupled with unidirectional airways, or multi-lane airways, step climbs could be further reduced to 1,000 ft, improving their effectiveness.

6.2 Speed Improvements

Many potential techniques are available to improve speed performance. As the analysis shows, speed improvements represent the largest area for fuel burn reduction on nearly all flights, making progress in this direction especially important. Additionally, speed changes to some extend should be easier to make. Changing speed does not require special requests from ATC, and unlike altitude changes they do not pose immediate monitoring difficulties for controllers. Ultimately, airlines and their pilots are less restricted by speed constraints. Unlike cruise climbs, which are not supported by current FMSs, speed optimums are stored in FMSs and can easily be commanded. This increased level of speed control, combined with the high fuel savings made possible by speed optimization, should make speed profile enhancements a natural strategy. Still, many barriers limit their effectiveness, but mitigation options for overcoming these barriers exist.

6.2.1 Custom Aircraft Speeds

The analysis shows that the optimal speed for a across various aircraft types is variable. This makes integration of these aircraft into the same NAS difficult at times, because aircraft can sometimes queue along single-lane airways. Faster aircraft conflict with slower ones ahead, causing nightmares for controllers of busy corridors. These conflicts result in either a diversion from the original flight path to avoid a conflict, or a change in speed for one of the aircraft. In the case of a speed change, this may alter the aircraft's efficiency significantly. Multi-lanes

enabled by RNP could allow for varying speed aircraft to occupy different lanes, allowing them to fly consistently at their speed optimum. At the least, ADS-B and RNP could allow for less drastic avoidance maneuvers along airways, allowing them to maintain speed while requiring only minor flight path deviations.

6.2.2 Speed Reduction Efforts

Perhaps the clearest lesson learned from viewing the aircraft speed profiles was that aircraft most always fly faster than their maximum SAR, or best fuel economy speed. This is due to a number of reasons but often related to economics, passenger satisfaction, and on-time performance. Time-sensitive costs will always dictate that the minimum cost airspeed will always be somewhat higher than the best economy speed. Because airlines are primarily concerned with cost, this force will be forever present. However, the data suggested that aircraft are operating at speeds even higher than typical CI-driven values. Therefore, reductions from these speeds would increase efficiency and reduce costs, and should be taken advantage of. These excessive speeds may be the product of poor planning or terminal delays which drive fast speeds to make up for lost time. Better operator planning and NextGen efforts to relieve congestion will likely help reduce speeds to more appropriate levels. For aircraft already operating at minimum cost speeds, only rising costs of fuel will serve to push airlines closer toward best fuel economy operations.

Alternatively, ATC could implement policies which favor slowing down. A less ambitious approach than custom airspeeds would be to slow faster aircraft down when presented with a traffic conflict along an airway. ATC would ultimately be commanding a slightly slower speed for most aircraft on the path, thus potentially pushing them closer to their best fuel economy speed. A potential downside to this idea is that if an aircraft is slowed beyond its best fuel economy speed, its fuel efficiency will decrease. An even simpler procedural change would be to give preference to slower aircraft when a conflict arises and path deviations are required. When

a faster aircraft approaches a slower one, the faster aircraft would be required to deviate to pass. This policy would encourage aircraft to slow down to avoid flight path deviations.

Chapter 7

Conclusion

The analysis results show promising potential for altitude and speed trajectory improvements. All flights showed some potential for fuel burn and emissions to be reduced via simple changes in the speed and altitude profiles, across a range of common aircraft and routes. The upper bound on fuel burn and emissions reduction for the flights and aircraft in this analysis is a 3.5% in cruise. This translates to 2.6% over the entire flight, given that the climb and descent phases were not affected. That result represents an annual domestic fuel burn reduction of 300 million gallons and CO_2 reduction of 3.2 million tons if every flight's speed and altitude trajectory were optimized.

While completely optimal speed and altitude trajectories are ideal, many barriers currently prevent their implementation. Cruise climbs are inhibited by controllers' limited ability to maintain safe aircraft separation, so cruise altitudes must be level except for periodic climbs and descents to other levels. This practical restriction makes step climbs a much more attractive solution in the near term. Step climbs of 2,000 ft are currently possible in today's RVSM environment, and 1.02% reduction in cruise fuel burn was seen in the set of flights analyzed when these step climbs were implemented. System-wide, implementing these operations would save 88 million gallons of fuel and cut CO_2 output by 950 thousand tons. Implementing smaller 1,000 ft step climbs increases the cruise savings slightly to 1.12%, and continuous cruise climbs

would bring this value up to 1.5%, or 130 million gallons. The value of enabling cruise climbs in the system is unclear given the challenges they pose, however taking advantage of the current 2,000 ft step climb capability is an obvious step in the right direction. Longer flights, over 1,000 nm, are most susceptible to increased fuel burn if step climbs are not utilized. In general, flights longer than this distance burn enough fuel to warrant a 2,000 ft altitude step.

Speed improvements show the highest potential for reducing fuel burn and emissions. Optimizing speed alone resulted in a 2.4% reduction in cruise fuel burn, nearly double that of the altitude-only optimization potential. However, implementation of best-economy cruise speeds is implausible. Changing winds call for varying optimal speeds, and these changes are likely to cause traffic conflicts on busy airways, resulting in path or altitude deviations to allow aircraft to pass. More importantly, operators will inevitably select speeds to minimize cost using a cost index (CI); the result is flying faster than the best fuel economy speed due to the consideration of time-sensitive costs. The data showed that in addition to most flights operating faster than the best fuel economy setting, many flights also operated faster than the LRC speed as well. LRC is thus a conservative selection as a speed reduction target. The results in this analysis showed that if all aircraft were to implement a constant LRC airspeed, fuel burn could be reduced 1.61% overall, saving 138 million gallons of fuel. Operating at lower CI settings would result in even lower speeds and higher efficiency. The implementation of this simple operational speed reduction offers to have a larger effect than any of the altitude improvements alone, making this a seemingly smart and easy policy to adopt.

Newer aircraft models like the Boeing 737, flown by leaner low-cost carriers showed the least available potential, due to the efficient operating policies already put in place by those carriers. Alternatively, legacy carriers flying aircraft like the 757 and MD82 showed the most potential for improvement from trajectory enhancements, especially those related to speed. This distinction illustrates that carriers' operating procedures can have a significant effect on

performance, and that deficiencies in the system are not wholly responsible for operational inefficiencies. Alternatively, even the best showed performance potential, indicating that the entire system has room for improvement.

A number of potential opportunities are available to increase the efficiency of commercial aircraft now and in the future. Unlike aircraft modifications, which are often expensive take years to migrate into the system, operational enhancements promise the ability to increase performance in the near term with no special hardware. Simple changes in the way aircraft speed and altitude are chosen can clearly have a system-wide impact on efficiency. As NextGen unravels and enables better trajectory control, system efficiency can be increased further. Operators, pilots, controllers, and regulators alike should embrace the fuel burn and emissions reduction potential available in today's NAS, and make steps toward taking full advantage it.

Appendices

Appendix A. Aircraft SAR Sensitivity to Speed and Altitude

Each contour line represents 1% reduction in SAR from the optimal point. Numbers are included on the contour lines to mark the percentage reduction in SAR.





Appendix B. Aircraft Optimal Altitude Versus Weight

Aircraft were run through a long range detailed mission profile in Piano-X and given free rein on altitude selection. The resulting altitudes and weights were recorded as the aircraft step climbed. These values were recorded at the start of each new step altitude, immediately after the step climb was complete.





Bibliography

- BTS. (2010, April 28). *Airline Fuel Cost and Consumption (U.S. Carriers All)*. Retrieved from RITA BTS Transtats: http://www.transtats.bts.gov/fuel.asp
- EPA. (2007, April 15). *Inventory of US Greehouse Gas Emissions and Sinks: 1990-2005.* Retrieved from Environmental Protection Agency: http://www.epa.gov/climatechange/emissions/downloads06/07CR.pdf
- FAA. (2009). *Enhanced Traffic Management System (ETMS)*. Retrieved from Federal Avation Administration: http://hf.tc.faa.gov/projects/etms.htm
- Hirsch, R. L. (2005, February). *Peaking of World Oil Production: Impacts, Mitigation, and Risk Management.* Retrieved from National Energy Technology Laboratory: http://www.netl.doe.gov/publications/others/pdf/Oil_Peaking_NETL.pdf
- Histon, J. M., & Hansman, R. J. (2008). *Mitigating Complexity in Air Traffic Control: The Role* of Structure-Based Abstractions. Doctoral Thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, Cambridge.
- IATA. (2010, December). *IATA*. Retrieved from Fact Sheet: Carbon Neutral Growth: http://www.iata.org/pressroom/facts_figures/fact_sheets/pages/carbon-neutral.aspx
- Kar, R., Bonnefoy, P. A., & Hansman, R. J. (2010). Dynamics of Implementation of Mitigating Measures to Reduce CO2 Emissions from Commercial Aviation. Cambridge: Massachusetts Institute of Technology.
- Kingsley-Jones, M. (2009, April 22). 6,000 and counting for Boeing's popular little twinjet. Retrieved from Flight International: http://www.flightglobal.com/articles/2009/04/22/325472/pictures-6000-andcounting-for-boeings-popular-little-twinjet.html
- NOAA. (2010, November 3). *About NOMADS*. Retrieved from NOAA Satellite and Information Service: http://nomads.ncdc.noaa.gov/nomads.php?name=about
- NOAA. (2010, November 19). *NCEP North American Regional Reanalysis*. Retrieved from NOAA Satellite and Information Service: http://nomads.ncdc.noaa.gov/docs/ncep-northamerica-1.pdf
- Roberson, W., Root, R., & Adams, D. (2007, September). *Fuel Conservation Strategies: Cruise Flight*. Retrieved from Boeing AERO Magazine: http://www.boeing.com/commercial/aeromagazine/articles/qtr_4_07/article_05_3.ht ml