Benefits of Continuous Descent Operations in High-Density Terminal Airspace Under Scheduling Constraints

John E. Robinson, III*
NASA Ames Research Center, Moffett Field, California, 94035

Maryam Kamgarpour†
University of California Berkeley, Berkeley, California, 94720

This study estimates the potential benefits of continuous descents for more than 480,000 flights to 25 major airports in the National Airspace System. While reduced fuel consumption and greenhouse gas emissions are expected for these procedures, the benefits during periods of congestion are not well understood. To address this gap, baseline trajectories were constructed from flight plan and track data for flights arriving at 8 busy terminal areas. Two types of continuous descent trajectories were modeled. One enforced a constant distance-to-fly constraint to simulate uncongested operations. The other enforced a constant time-to-fly constraint to simulate congested operations. Potential fuel savings were calculated for different continuous descent scenarios. Analysis of the distance-constrained trajectories showed that fuel savings was distributed unevenly among the flights. The estimated savings was less than 25 kg for over 45% of the flights, and less than 100 kg for over 87% of the flights. The time-constrained trajectories showed 70-85% less potential savings than the distance-constrained trajectories. Prioritization of the improvements necessary to execute continuous descents during periods of congestion must rely upon analysis of a sufficient sample of operations, representative of many days, aircraft types, and traffic demand levels.

I. Introduction

THE Next Generation Air Transportation System (NextGen) plan mandates the development of advanced air traffic management technologies and procedures to accommodate safely, efficiently, and reliably a significant increase in traffic demand in the already congested terminal environment.[1] A concept of operations for NextGen terminal airspace, referred to as Super-Density Operations, envisions the use of advanced ground and flight deck automation, efficient RNAV/RNP routes, optimized vertical profiles, and delegated interval management to maintain high airspace and airport utilization in all weather conditions.[2] Lately, the costs associated with inefficient flight paths, particularly vertical profiles, has increased because of the rising cost of fuel and the overall financial weakness of the airline industry.

One method for increasing the vertical profile efficiency is a Continuous Descent Approach (CDA). CDAs are arrival procedures in which the aircraft descends continuously from cruise to landing with engines at or near idle. CDAs contrast against today's typical "dive and drive" procedures in which aircraft fly powered constant-altitude segments, referred to as level segments, at intermediate altitudes after their initial descent from cruise altitude. Flying these level segments, which may be chosen on the basis of local traffic, schedule constraints, and weather, generally increases fuel burn, greenhouse gas emissions, and noise pollution. For this reason, the CDA trajectory has also been referred to as a "green trajectory."

Operationally, two types of continuous descents have been implemented: Optimized Profile Descents (OPDs) and Tailored Arrivals (TAs). OPDs are improved Standard Terminal Arrival Routes (STARs) which allow aircraft to perform continuous descents until interrupted by air traffic controllers. OPDs are published, thus statically defined, arrival procedures designed with altitude restrictions that accommodate a wide-variety of aircraft types across a

* Aerospace Engineer, Aerospace High Density Operations Branch, Mail Stop 210-6, AIAA Senior Member.
† Graduate Student Researcher, Department of Mechanical Engineering, Mail Code 1740, AIAA Member.

range of expected weather conditions. OPDs have been used operationally in the United States at LAX[3], ATL[4], SDF[5], as well as, LAS, MIA, PHX, and SAN. TAs are aircraft-specific, thus dynamically defined, 4-D trajectories generated by advanced air traffic management automation that account for local traffic, schedule constraints, weather, and the aircraft's individual capabilities. TAs have been used operationally in the United States at SFO[6], LAX, and MIA.

Although the potential benefits of continuous descents on an individual flight basis seem clear, the execution of continuous descents has proven difficult during periods of traffic congestion. One challenge is that in order for a continuous descent to be uninterrupted, traffic separation demands and airspace restrictions cannot be allowed to interfere with the flight's descent profile. Hence, these optimized descent profiles have been feasible only during periods of low to moderate traffic demand.

Several studies have attempted to quantify the potential benefits of continuous descents through analyses of fuel burn, greenhouse gas emissions and noise reduction using limited sample sizes (in terms of the number of flights, and aircraft types studied).[6-11] There have however, been fewer analogous studies for large samples of flights to individual airports.[12-14] Another approach has been to extrapolate a small number of flights to the entire NAS.[15] The major contribution of this study is the estimation of the potential benefits of continuous descents for a large sample of flights representing many airports and aircraft types during both uncongested and congested periods of operations. The results of these analyses are reported in terms of probability distributions and the variation of those distributions for different airports, aircraft weight-classes, and operational scenarios. The paper is organized as follows. Some previous work that estimated the potential benefits of continuous descents is reviewed in Sec. II. The objectives of the current study, influenced by the scope and limitations of the previous work, are described in Sec. III. In Sec. IV, the analysis methodology is described in detail. The results are presented in Sec. V. Their implications, and relationship to previous findings are discussed in Sec. VI. Finally, in Sec. VII, conclusions are made regarding the implementation of continuous descents in congested airspace.

II. Background

Originally, CDAs referred to a truly continuous idle thrust descent from cruise to touchdown. As a result of operation at idle thrust, these optimal descent profiles are realizable only during periods of low traffic demand. In this paper, the term, continuous descent, will refer more generally to any optimized vertical profile that has fewer or no unnecessary powered level segments at intermediate altitudes.

Several studies have estimated the potential fuel savings of specific optimized arrival procedures using sophisticated fuel flow models and actual flight test data for a specific, often limited, sampling of flights. Coppenbarger, based upon the use of Boeing's proprietary database and aircraft performance analysis software, estimated the benefits of two Oceanic TA trajectories for varying levels of traffic demand at SFO.[6] Similarly, Tong estimated the benefits of a new arrival procedure for dual runway operations at IAH.[7] Others have analyzed actual flight test data from operational trials. Wat used flight recorder data to estimate the benefits of more than 150 continuous descent operations at Schipol Airport (EHAM).[8] Similarly, Clarke used flight recorder data to estimate the benefits of more than 120 continuous descent operations at SDF.[5] Reynolds used flight recorder data to separate the path, vertical and speed inefficiencies of current operations using approximately 1800 European Airbus A320 flights.[9] Reynolds' analysis allows changes to operational procedures to target the most critical sources of inefficiency.

All of the studies cited in the previous paragraph share a number of limitations. One is small sample size namely, the number of flights, the mix of aircraft types and the traffic demand levels. Another is the predominant focus on heavy aircraft, whose potential fuel savings are substantially higher per aircraft than for other aircraft types. Yet another limitation is that the operational trials of advanced arrival procedures were restricted to periods of low to moderate traffic demand when aircraft interactions were less likely. Among these studies, only Reynolds' work distinguished between the sources of the fuel savings specific to the improved operational procedures, such as shorter horizontal routes, optimized speed profiles, and optimized vertical profiles.

Assessments of potential fuel savings of continuous descents using simpler fuel flow models and today's traffic for individual airports have appeared in the literature. Sprong used EUROCONTROL's Base of Aircraft Data (BADA)[10] aircraft performance model to estimate the benefits of 10 AIRE demonstration flights at MIA using 16 days of traffic and 11 AIRE demonstration flights at ATL using 6 days of traffic.[11] Dinges used BADA to estimate the fuel savings of continuous descents using four days of LAX traffic.[12] More recently, Shresta used BADA to estimate the fuel savings of full and partial continuous descents in the presence of airspace constraints using an unspecified number of days of DEN traffic.[13] Compared to the previously mentioned studies, these studies have larger sample sizes and more realistic mixes of aircraft types. As a result, their predicted benefits are substantially less on average per aircraft. The technical approach taken in these studies has allowed isolation of the benefits of improved vertical profiles from other sources of fuel savings. However, each study was associated with a single airport, and the number of days of traffic was small.

Finally, a number of studies have estimated the potential fuel savings of continuous descents for the entire NAS. Melby used BADA to estimate the NAS-wide fuel savings of continuous descents using one day of traffic at the OEP 35 airports.[14] Similarly, Alcabin used an unspecified model to estimate the NAS-wide fuel savings of both continuous climbs and descents by extrapolating only 81 ORD flights and 61 DEN flights to the OEP 35 airports.[15] These studies were aimed at quantifying the potential NAS-wide benefits of continuous descent operations. However, these studies, like earlier ones conducted for individual airports, analyzed a small number of days for their constituent airports, and made no attempt to translate the results to other airport and runway configurations, traffic demand levels, or delay conditions.

III. Motivation

The principal goal that led to this study is to determine an operational concept for efficient descent procedures in congested airspace. This objective is pursued by:

- estimating the benefits of continuous descents for current and future traffic demand and airborne delay levels (i.e., How do the benefits scale in tomorrow's presumed more congested airspace?) and,
- investigating prioritization of airspace, procedure, and technology changes needed to achieve use of continuous descents during all traffic conditions (i.e., How widely and frequently should continuous descent procedures be used?).

Numerous studies provide operational concepts for the execution of continuous descents in the NAS. Although, it may seem that an ultimate goal for the arrival phase of flight is "continuous descents from cruise to touchdown all of the time for all flights," the costs and benefits of such operations have not been fully explored. Continuous descents, by their very nature, have tangible economic and environmental benefits, but the infrastructure and technologies needed to achieve these benefits depend upon the specific operational concept, and the associated costs can be substantial. Previous studies do not aggregate and report the benefits of continuous descents in a manner that helps prioritize the airspace and procedures that should be improved or the technologies that should be developed.

Furthermore, it is believed that the majority of potential fuel savings for continuous descents is associated with level segments in terminal airspace. While execution of level segments at lower altitudes will generally have the greatest potential fuel penalty, the technology needed to eliminate them will also generally be the most complex, and therefore costly. A systematic analysis of possible continuous descent scenarios using traffic samples covering a large number of operational conditions to different airports is needed to support informed investment decision-making.

Fuel savings are often reported as the absolute weight of fuel saved with no baseline for comparison, such as the total fuel consumption. Typically, benefits are described in terms of average fuel savings per flight or total fuel savings for all flights. This is appropriate when benefits are distributed among the flights evenly. However, when benefits for some subsets of flights are significantly greater than for others, the large standard deviation of the entire sample makes the mean uninformative. Therefore, greater details should be reported in the results than simply general statistics.

Lastly, previous studies have assumed that advanced scheduling capabilities will allow flights to be scheduled before, even well before, their initial top-of-descent. Ideally, flights could then use minor speed corrections to absorb any additional delay needed past top-of-descent. For this condition, the distance-to-fly of the route is essentially fixed, and the scenario can be formulated as a maximum range problem (i.e., minimize fuel burn per distance). However, consideration needs to be given to a scenario where flights must absorb delay in the form of level segments past their initial top-of-descent, such as during periods of high traffic demand. For this condition, time along the route is essentially fixed, and the scenario should be reformulated as a maximum endurance problem (i.e., minimize fuel burn per time).

Motivated by the research questions stated above and the scope of the previous research, our analysis methodology is designed to explore the potential benefits of continuous descents for various aircraft types and airports, subject to varying levels of traffic demands. We present the results of our study such that the inefficiencies of the current procedures at various altitudes for different airports can be compared. In addition, we estimate the amount of potential fuel savings for the cases in which delay needs to be absorbed in the form of level segments.

IV. Analysis Methodology

This research estimates the potential fuel savings of continuous descents for more than 480,000 flights. Arrival flights to multiple airports in multiple TRACONs were analyzed in order to encompass congested metroplex environments, like New York, as well as super-hub environments, like Atlanta. The traffic analysis was performed using NASA's Center/TRACON Automation System infrastructure in conjunction with ARTCC and TRACON flight plan and track data.[16] The Aviation Systems Division at NASA Ames Research Center has 24 hour-a-day, 7 day-a-week access to the following data:

- flight plan and track data for all 20 ARTCCs, as well as 27 TRACONs in the NAS,
- NOAA Rapid Update Cycle (RUC) weather forecasts, and
- NWS Meteorological Aviation Reports (METAR).

Fuel flow analyses were performed using BADA and several Aircraft Operating Manuals. The results were imported into an SQL data warehouse in order to investigate the macroscopic factors that strongly affect vertical flight path efficiency. The potential fuel savings were aggregated by airport, arrival procedure, landing runway, aircraft weight-class, and traffic demand level. More than 2000 CPU hours were spent processing the raw data on high-performance Apple™ computers. The volume of the resulting data warehouse exceeded 1 terabyte of raw, processed, and database-resident data.

A. Scope of Study

The vertical flight paths were analyzed for arrival flights into 8 TRACONs in the NAS, including all seven so-called Large, or Consolidated, TRACONs. The seven Large Level 12 TRACONs are Atlanta TRACON (A80), Chicago TRACON (C90), Dallas/Fort Worth TRACON (D10), New York TRACON (N90), Northern California TRACON (NCT), Potomac TRACON (PCT), and Southern California TRACON (SCT). The vertical flight paths for arrival flights into Denver TRACON (D01) were included for comparison to Ref. 13. Arrival flights to these eight TRACONs' OEP 35[17] airports, as well as to their major satellite airports, were analyzed. Table 1 summarizes the traffic statistics for these facilities.

Table 1. Summary of TRACONs and airports that were studied.

Name	ID	2009 Ops Count[1]	2009 Ops Rank[1]	OEP 35 Airports	Satellite Airports
Atlanta	A80	1,275,000	5	ATL	-
Chicago	C90	1,186,000	6	ORD MDW	-
Dallas/ Fort Worth	D10	1,066,000	7	DFW	DAL
Denver	D01	853,000	10	DEN	-
New York	N90	1,831,000	2	EWR JFK LGA	HPN TEB
Northern California	NCT	1,446,000	4	SFO	OAK SJC SMF
Potomac	PCT	1,490,000	3	BWI DCA IAD	RIC
Southern California	SCT	1,972,000	1	LAX SAN	BUR LGB ONT SNA

[1] Source: FAA Administrator's Fact Book, March 2010

For each TRACON, arrival flights were analyzed for between thirty and sixty 24-hour traffic samples. These daily traffic samples were distributed between February 2010 and early May 2010. The traffic samples covered the entire day starting at 0800Z (between 0100 and 0400 local time) on the day of interest and ending at 0800Z on the next day. Table 2 summarizes the days and the number of operations captured by these traffic samples. The traffic samples chosen for analysis were required to have uninterrupted ARTCC and TRACON track data during the busiest period of the day, specifically between 0600 local time and 2200 local time. Uninterrupted data was available for all hours of most days.

The proportions of flights by weight-class[18] were 81% large, 8% heavy, 7% Boeing 757, and 5% small, plus fewer than 150 individual super-heavy flights. Regional, business and micro jets constituted more than 35% of the jet traffic. The proportions of flights by engine type were 93% jets, 6% turboprops, and 1% pistons. The traffic

Table 2. Summary of 24-hour daily traffic samples that were studied.

Name	ID	Feb. 2010		Mar. 2010		Apr. 2010		May 2010		Estimated Share of 2010 Ops[1]
		Days	Arrival Ops	Days	Arrival Ops	Days	Arrival Ops	Days	Arrival Ops	
Atlanta	A80	3	3,347	11	15,520	13	18,147	5	6,456	9%
Chicago	C90	13	17,733	13	19,613	17	26,420	5	7,649	13%
Dallas/ Ft. Worth	D10	19	19,467	16	19,449	21	23,328	4	4,522	16%
Denver	D01	4	3,312	14	12,429	15	14,309	0	0	10%
New York	N90	14	22,814	17	32,296	23	42,964	5	10,592	14%
Northern California	NCT	16	15,498	21	21,351	8	7,251	4	3,997	10%
Potomac	PCT	2	1,880	10	12,262	20	22,733	4	4,059	8%
Southern California	SCT	10	14,951	19	27,913	17	25,409	6	8,985	12%

[1] Source: www.airnav.com, June 2010.

samples span approximately one-quarter of the 2010 calendar year and more than 15% of the expected 2010 arrival operations to the 23 airports listed in Table 1.

These traffic samples were not explicitly chosen to represent all possible airspace and airport configurations or weather conditions. A summary examination of the airport configurations in use during each traffic sample suggests that many of the most common airport configurations, if not all of them, are included in the analysis. A subsequent and more complete accounting of the airport configurations is desired, since less-common airport configurations may exhibit larger inefficiencies caused by their low frequency of occurrence.

B. Analysis Using Flight Plan And Radar Tracks

Analysis of the potential fuel savings for continuous descents followed these steps: characterization of the operational scenario (i.e., landing runway, landing time, flight distance, etc.), identification of level segments, and calculation of potential fuel savings associated with moving level segments to higher altitudes. The remainder of this section will discuss these steps in detail.

1. Landing Runway Determination

Each flight's landing runway was determined using simple geometric heuristics. An imaginary rectangular bounding box was constructed along the runway's final approach course. Longitudinally, it extended 10 nmi from the runway threshold. Laterally, it extended 0.25 nmi away from the runway centerline or bisected the centerlines of parallel runways. An aircraft was assigned to a particular runway if its three most recent track positions were inside the bounding box, its course was within 1 degree of the runway's localizer, its turn rate was less than 0.5 degrees per second, and it was not more than 500 feet above the runway's glide slope. In order to achieve reliable runway determination results and to align the track positions with the appropriate physical landmarks, the numerical precision of the latitude and longitude coordinates of each TRACON's radar sensors needed to be 0.5 arc-seconds, and the numeric precision of their magnetic variances needed to be 0.1 degrees.

Visual inspection of all traffic samples showed that the simple geometric heuristic correctly identified more than 99.5% of the landing runways. Each day, most airports had fewer than 10 flights that were assigned an incorrect landing runway or were not assigned a runway. The simple geometric heuristics did not perform well for two airports, DCA and LGA. These airports' curved visual approaches and occasional circle-to-land operations made the simple geometric heuristics insufficient, and caused the number of flights without a correct landing runway to exceed 20%. These airports are excluded from those analyses that require accurate landing runway information. Figure 1 shows flights arriving to ATL, DFW, ORD, and SFO on May 4, 2010. Each flight is color-coded by its landing runway. The track position accuracy was sufficient to identify the separate traffic flows to even the most closely-spaced parallel runways.

2. Runway Threshold Crossing Time Determination

The flight's landing runway and runway threshold crossing time were used to calculate the airport and runway arrival rates. Each flight's runway threshold crossing time was calculated by interpolating between two track positions that straddled the runway threshold or by extrapolating the flight's last track position to the runway

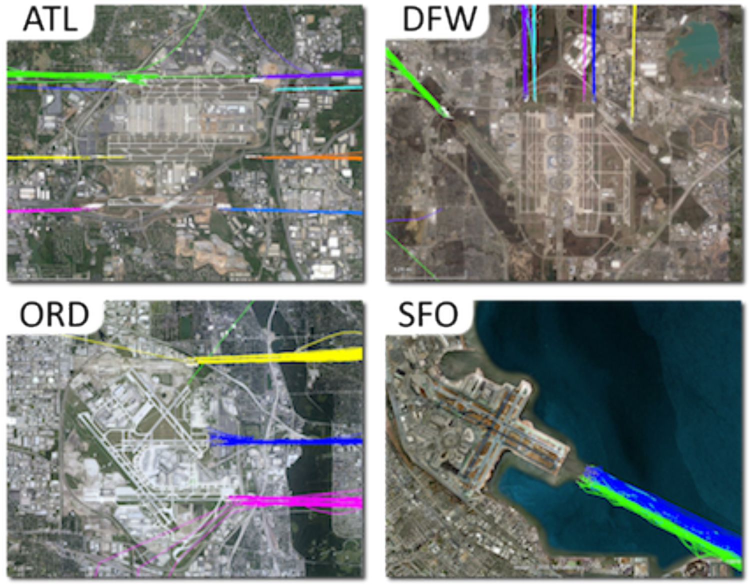

Figure 1. Examples of ATL, DFW, ORD and SFO arrivals color-coded by landing runway for May 4, 2010.

threshold at constant ground speed. A limited comparison of DFW ASDE-X[19] airport surveillance radar data (1 second update period) to D10 ARTS[20] radar data (4.8 second update period) found that the error in the predicted runway threshold crossing time was less than five seconds. Analysis of all traffic samples showed that, with the exception of DCA and LGA for the reasons described above, runway threshold crossing times were predicted for more than 99% of the arrival flights.

3. Identification of Level Segments

For each flight's track history, a piecewise linear vertical profile, consisting of ascent, descent and level segments, was computed. The level segments in each flight's vertical profile were determined using the set of pattern-matching heuristics shown in Table 3. The level segments were considered to start and end at the first and last track positions that satisfied the heuristics. The ascent and descent segments were a natural byproduct of connecting the endpoints of the level segments. The heuristics were purposely designed to find level segments at standard flight levels (e.g., 5000 feet MSL) more easily than level segments at non-standard altitudes (e.g., 5300 feet MSL).

Table 3. Pattern-matching heuristics used for the identification of level segments from track data. The rules are read from left to right.

If flight is	0-100 feet above	flight level	for longer than	15 seconds	then, flight is level at that	flight level
	0-100 feet below	flight level		15 seconds		flight level
	constant at	any altitude		30 seconds		altitude
	within 100 feet of	flight level		40 seconds		flight level
	within 100 feet of	any altitude		50 seconds		altitude

Visual inspection of five daily traffic samples of SCT arrival flights showed that the pattern matching heuristics correctly identified all of the visually recognizable level segments. The coarse 100-foot resolution of the Mode C altitude reports combined with the 4.8-second update period of the ARTS radar made it difficult to recognize level segments whose durations are less than 15 seconds. Figure 2 shows the vertical profile of a flight whose level and descent segments have been identified.

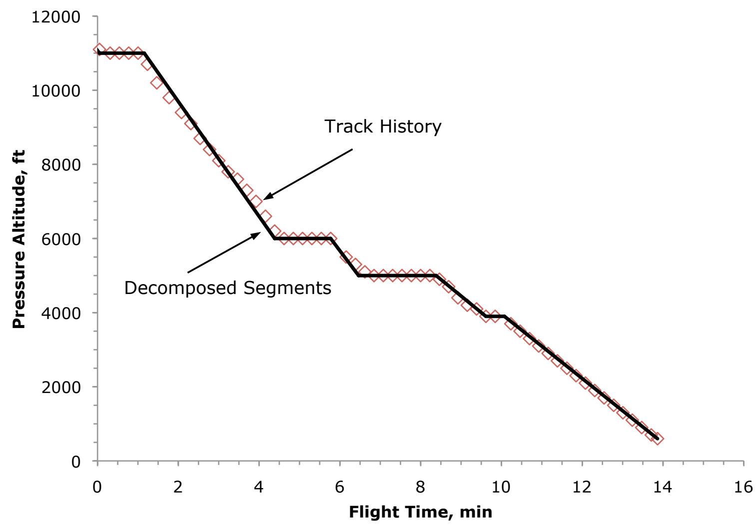

Figure 2. Example of vertical profile decomposition for a flight having three intermediate level segments.

4. Calculation of Level Segment Fuel and Flight Time Savings

BADA was used to estimate the fuel and flight time savings associated with replacing particular level segments at lower altitudes by level segments at higher altitudes. Only the level segments were analyzed, since the fuel consumption and duration of a series of interrupted descent segments was assumed to be similar to that of a single continuous descent segment. Each flight's baseline trajectory was used to calculate two new theoretical trajectories: one constrained by distance-to-fly and one constrained by time-to-fly. The calculations of the flight's potential fuel savings and flight time reduction, subject to these constraints, are described below.

Distance-Constrained Trajectory: The distance-to-fly of the trajectory is prescribed, whereas the time-to-fly is allowed to vary. BADA's nominal aircraft type-specific cruise speed profiles, rather than the observed speeds, were used to calculate the fuel flow rate of the level segments. This approach ensured that statistically similar speeds were used at all altitudes. This scenario corresponds to a fixed distance-to-fly matching the baseline trajectory's horizontal path. It is representative of low to moderate traffic demand when the time-to-fly is not constrained by traffic metering or separation. It is also assumed to represent the upper bound of expected fuel and flight time savings. Equations (1) and (2) describe the fuel savings and flight time decrease associated with moving all level segments to a higher altitude, respectively.

$$\Delta TTF = \sum_{i=0}^{N} d_i \cdot \left(\frac{1}{V_{TAS}(h_i)} - \frac{1}{V_{TAS}(h'_i)}\right) \tag{1}$$

$$\Delta FC = \sum_{i=0}^{N} d_i \cdot \left(\frac{FF(h_i)}{V_{TAS}(h_i)} - \frac{FF(h'_i)}{V_{TAS}(h'_i)}\right) \tag{2}$$

where,
ΔTTF is the time-to-fly change associated with moving all level segments,
ΔFC is the fuel consumption change associated with moving all level segments,
d_i is the original length of the i'th level segment,
h_i is the original altitude of the i'th level segment,
h'_i is the new altitude of the i'th level segment,
$V_{TAS}(h)$ is the nominal cruise speed of the i'th level segment at a particular altitude, and
$FF(h)$ is the cruise fuel flow rate of the i'th level segment at a particular altitude.

Figure 3 illustrates the notional distance-constrained translation of a particular level segment to one at higher altitude and faster speed. The length of each level segment remains unchanged in the translation (Fig. 3a), so the durations of the level segments, and thus the time-to-fly of the overall route, are decreased (Fig. 3b).

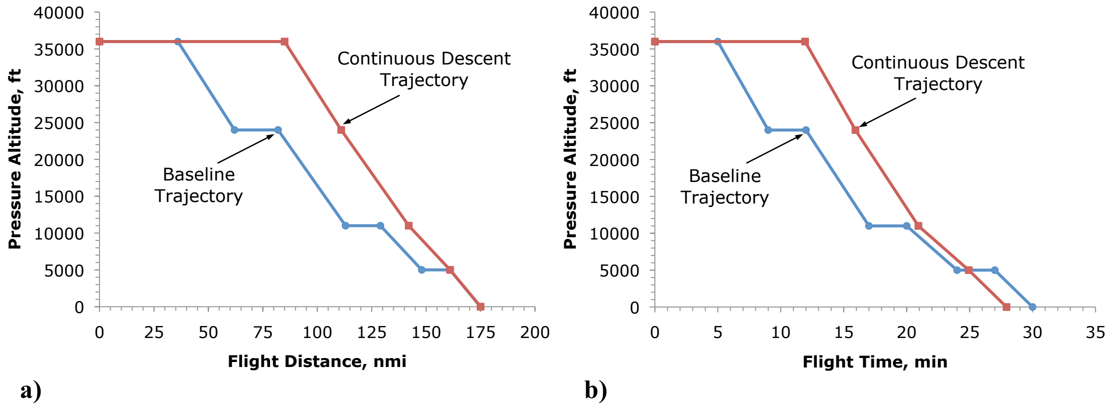

Figure 3. Illustration of distance-constrained translation of baseline trajectory to continuous descent trajectory.

Time-Constrained Trajectory: The time-to-fly of the trajectory is prescribed, whereas the distance-to-fly is allowed to vary. A nominal aircraft type-specific holding speed at each level segment's original altitude, rather than the observed speeds, was used to calculate the fuel flow rate of the level segment at its new altitudes. This approach ensured that statistically similar speeds were used at all altitudes. This scenario corresponds to a fixed time-to-fly matching the baseline trajectory's time of arrival. It is representative of moderate to high demand when the time-to-fly is constrained by traffic metering or separation. It is also assumed to represent the lower bound of expected fuel and flight time savings. Equations (3) and (4) describe the fuel savings and distance-to-fly increase associated with moving all level segments to a higher altitude, respectively.

$$\Delta DTF = \sum_{i=0}^{N} t_i \cdot (V_{TAS}(h_i') - V_{TAS}(h_i)) \tag{3}$$

$$\Delta FC = \sum_{i=0}^{N} t_i \cdot \left(FF(h_i) - FF(h_i')\right) \tag{4}$$

where,
ΔDTF is the distance-to-fly change associated with moving all level segments,
ΔFC is the fuel consumption change associated with moving all level segments,
d_i is the length of the *i*'th level segment,
t_i is the duration of the *i*'th level segment,
h_i is the original altitude of the *i*'th level segment,
h_i' is the new altitude of the *i*'th level segment,
$V_{TAS}(h)$ is the nominal holding speed of the *i*'th level segment at a particular altitude, and
$FF(h)$ is the cruise fuel flow rate of the *i*'th level segment at a particular altitude.

Figure 4 illustrates the notional time-constrained translation of a particular level segment to one at a higher altitude and faster speed. The duration of each level segment remains unchanged in the translation (Fig. 4b), so the lengths of the level segments, and thus the distance-to-fly of the overall route, are increased (Fig. 4a). This model of the time constrained is deliberately pessimistic, since it does not use descent speed changes to achieve the required delay.

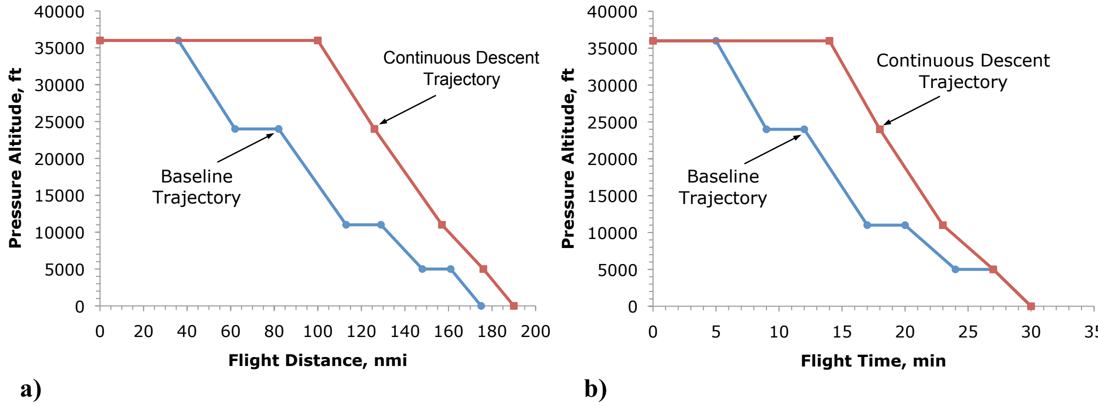

Figure 4. Illustration of time-constrained translation of baseline trajectory to continuous descent trajectory.

5. Aggregation of Level Segment Fuel and Time Savings

The individual fuel savings and flight time savings values for each level segment were combined for different concepts of full and partial continuous descents. Four key altitudes were defined as the possible boundaries between traditional descent procedures and continuous descent procedures. These altitudes were the flight's final cruise altitude (varied by flight), a nominal en route high sector to low sector transition altitude (FL240), a nominal TRACON ceiling (17,000 feet MSL), and a nominal arrival metering fix altitude (11,000 feet MSL). Two scenarios were examined for each altitude: a full continuous descent scenario that allowed all level segments below the key altitude to be moved to the flight's final cruise altitude and a partial continuous descent scenario that allowed all level segments below the key altitude to only be moved to that altitude and no higher. The resulting seven scenarios can be compared to determine the relative fuel savings of continuous descents during various stages of the descent profile.

C. Adjusted Fuel Flow Model

Previous research can be divided into two categories: studies with smaller sample sizes and higher fidelity fuel flow models, and studies with larger sample sizes and lower fidelity fuel flow models. The use of higher fidelity fuel flow models is hindered for large sample sizes by the lack of precise intent information, especially the speed profile and winds aloft. Another disadvantage of these models is that they are often difficult to integrate into large-scale analyses, and are subject to Intellectual Property and Company Proprietary Data restrictions. This study made three adjustments to the BADA (version 3.7.2) fuel flow model. These adjustments were intended not to develop a precise fuel flow model but to correct several gross errors, and to show the sensitivity of potential fuel savings to the fuel flow model itself. These adjustments included the following:

- compensation for the under-prediction of cruise fuel flow rate with respect to aircraft operating manuals,
- reduction of the nominal airspeed to match typical speed profiles in congested terminal airspace, and
- compensation for aircraft configurations using flaps at lower altitudes and slower airspeeds.

The magnitudes of these adjustments varied with the altitude of the level segment and aircraft type.

1. Adjustment of the BADA Cruise Fuel Flow Rates at Various Altitudes

A comparison of the fuel flow predicted by BADA and the fuel flow values listed in the aircraft operating manuals of several Boeing and Airbus commercial jet aircraft was made in order to quantify the accuracy of the BADA model. Figure 5 shows the difference between the BADA predictions and the published manufacturer values. The data are shown for six different aircraft types in the FLAPS UP configuration at their recommended holding speeds at each altitude. At typical en route altitudes (above FL240), the fuel flow rates predicted by BADA are within 5% of the published manufacturer values. This is consistent with earlier work that found the mean fuel consumption error in cruise to be less than 5%.[21] However, at terminal altitudes (below 11,000 feet MSL), the BADA fuel flow estimates are under-predicted by 15-25% when compared to the manufacturer's data. Recently, Senzig reported that comparisons of BADA fuel flow estimates and the airline's reported fuel consumption during a climb to 3000 feet above field elevation (AFE) have a mean under-prediction error of 22%.[22] One possible source of this error is the reduction of engine efficiency at lower altitudes and slower true airspeeds; this reduction is not modeled by BADA. In order to correct this modeling error, the fuel flow predicted by BADA was multiplied by a coefficient that varies linearly from 125% at sea level to 100% at FL350. This form for the coefficient was chosen for two reasons. First, the rate of change of the under-prediction error with altitude was similar for the common types of jet aircraft. Second, the coefficient's magnitude was similar to the Boeing 737-400, and Roach showed that the aggregate fuel burn profiles of today's operations to busy airports are reasonably well modeled by a Boeing 737-class aircraft.[23]

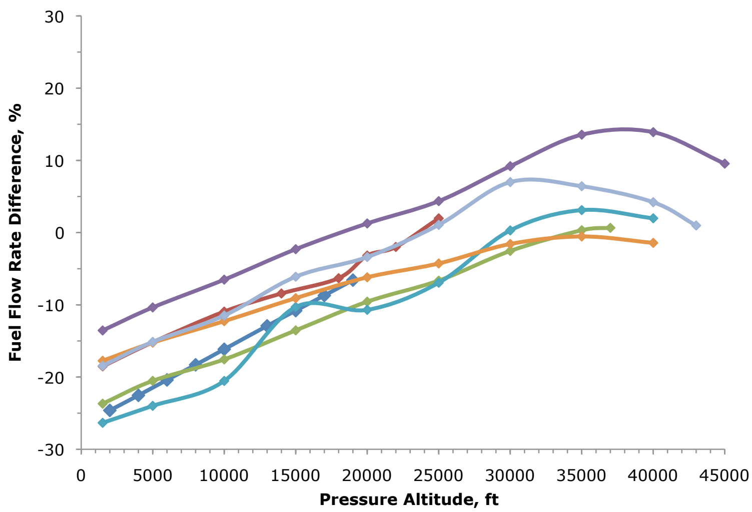

Figure 5. Comparison of BADA fuel flow predictions and published values from the aircraft operating manuals of several Boeing and Airbus aircraft (FLAPS UP configurations at recommended holding speeds).

2. Adjustment of the BADA Cruise Airspeed Profile in Terminal Airspace

Due to lack of access to information on the aircrafts' true airspeed and for modeling simplicity, the median ground speed was assumed to be a proxy for the median true airspeed. This approximation is reasonable given the large number of flights across a wide geographic and temporal range. Figure 6 shows the juxtaposition of the median nominal cruise speed profile assumed by BADA and ground speeds observed for all level segments. The blue and dark red lines show the nominal cruise speed profile and the median ground speeds observed for all level segments (approximately 1,400,000 level segments and more than 480,000 flights), respectively. The light red lines show the 25^{th} percentile (slower) and 75^{th} percentile (faster) ground speeds observed for these level segments. The large variance is caused by the combination of unmodeled winds aloft, variation of pilot behavior and aircraft performance, unaccounted for controller commanded speeds, and the inherent inaccuracy of the track data. Conversion of the observed ground speeds to their corresponding true airspeed values using RUC weather forecasts did not show markedly reduced variance. At terminal altitudes (below 5000 feet AFE), the nominal cruise speed profile diverges from the median ground speed of the level segments. In order to account for this speed inaccuracy, the BADA cruise speed profile was multiplied by a coefficient that varied linearly from 65% at 0 feet AFE to 100% at 4000 feet AFE. The black line shows the adjusted cruise speed profile. This adjusted speed was used in Eqs. (1) through (4). Only the low altitude portion of the cruise speed profile was adjusted. Specifically, the low altitude portion of the cruise speed profile was adjusted to be statistically similar to the remainder of the profile, i.e., it is near the 75^{th} percentile bound rather than near the median. It is noted that the correlation of landing direction and wind direction would naturally imply that the observed ground speeds should be less than the cruise speed profile along the final approach course.

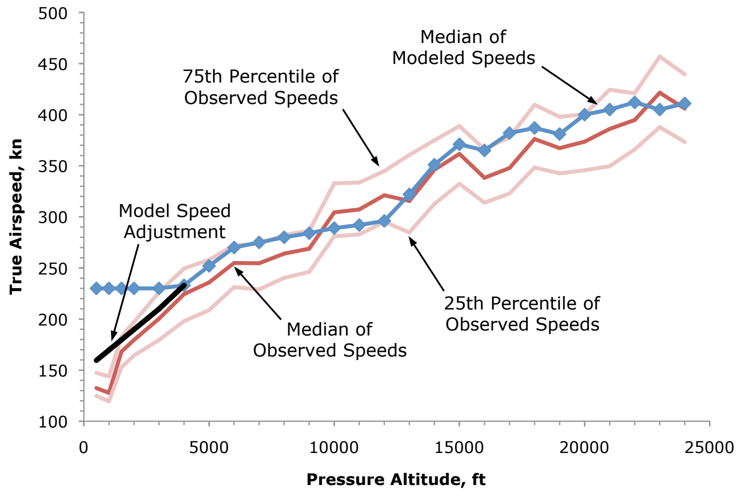

Figure 6. Comparison of BADA nominal cruise speed profiles and observed ground speeds on level segments.

3. Adjustment of the BADA Fuel Flow Estimates for Flaps

The use of high-lift devices (i.e., flaps) can dramatically increase the fuel flow rate for aircraft at low altitudes and speeds. For this study, three aircraft configurations were modeled: FLAPS UP, FLAPS 1°, and FLAPS 5°. The assumed minimum maneuvering speeds were chosen as follows:

• Heavy jets: 200 KIAS in the FLAPS UP configuration and 180 KIAS in the FLAPS 1 configuration
• Large jets: 190 KIAS in the FLAPS UP configuration and 170 KIAS in the FLAPS 1 configuration

Figure 7 shows the fraction of level segments whose observed ground speeds are slower than the aircraft's prescribed minimum maneuvering speeds for FLAPS UP, FLAPS 1 and FLAPS 5 aircraft configurations. The fuel flow penalties prescribed to the FLAPS 1 and FLAPS 5 configurations were 15 and 35%, respectively. These numbers were derived from the aircraft operating manuals of several Boeing and Airbus commercial jet aircraft. A Boeing 747-400 FAA-certified Level D simulator[24], operated by the Crew-Vehicle Systems Research Facility at NASA Ames Research Center, was used to independently validate these fuel flow penalties for a series of indicated airspeed values at terminal altitudes. The probability distributions of the observed ground speeds of level segments at each particular altitude and the fuel flow penalties of each aircraft configuration were used to derive an aggregate fuel flow penalty for all level segments as a function of their altitude. The fuel flow was multiplied by a coefficient that varied linearly from 140% at 0 feet AFE to 100% at 5000 feet AFE. Qualitatively, this adjustment corresponds to the assumption that one-quarter of the flights would be in a FLAPS 1 or FLAPS 5 configuration at 4000 feet AFE, two-thirds of the flights would be in those configurations at 3000 feet AFE, and nearly all of the flights would be in them below 3000 feet AFE. Even higher penalties for low altitude flight can arise from specific company procedures mandating the use of flaps at specific spatial locations along the approach path. An example of such a procedure is "Pilot Flying will call 'FLAPS 1, SPEED' not later than 7NM prior to the final approach fix on a straight in approach".[25]

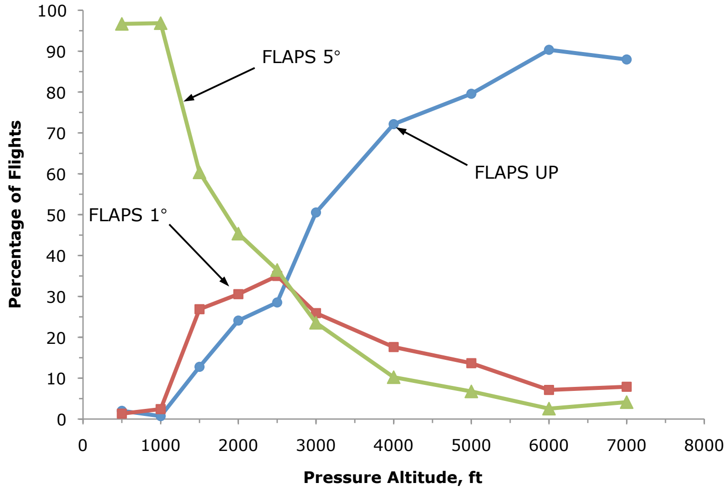

Figure 7. Percentage of level segments associated with different aircraft configurations.

4. Comparison of original and adjusted fuel flow predictions

The Boeing 737-800 is a modern single-aisle, narrow-body, twin-engine jet aircraft, and it is representative of a large share of today's operations in congested airspace. Figure 8 shows the fuel flow predicted by BADA for a Boeing 737-800 (59,800 kg, ISA, FLAPS UP) using the adjustments to the fuel flow rate and cruise airspeed profile described in the previous sections. Figures 8a and 8b show the fuel flow rate per unit distance and fuel flow rate per unit time, respectively. The original BADA values are shown in blue; the adjusted BADA values are shown in red. The adjustments to the fuel flow rate and cruise airspeed profiles increase the predicted fuel flow by 20 to 40% for level segments at terminal altitudes. For the nominal cruise speed profile, both models agree that flying higher minimizes fuel flow rate per unit distance. Fuel flow rate per unit time for this profile, however, has less variation with altitude. In Fig. 8b, the original BADA model suggests that flying lower minimizes the fuel flow rate per unit time, whereas the adjusted BADA model shows that there is a modest variation (with altitude) in the opposite direction at low altitudes.

These results illustrate the need for accurate fuel flow models across the entire arrival domain. Small corrections to the model have made a qualitative difference in the fuel flow per unit time. Such models are critical in determining how to manage efficient vertical profiles in the presence of time constraints. Inaccurate fuel flow estimates can lead to principally erroneous conclusions. For example, low fidelity models might suggest that delay, when it must be absorbed in the form of level segments, is best absorbed at lower altitudes. Worse yet, the application of fuel flow per unit distance might indicate that the fuel savings of efficient vertical profiles grows superlinearly with traffic demand and its corresponding arrival delay.

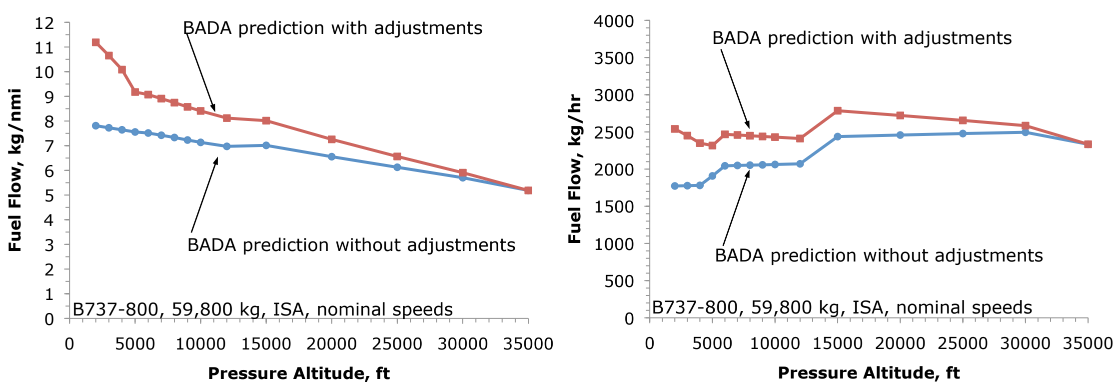

a) b)

Figure 8. Estimates of fuel flow per unit distance and fuel flow per unit time for BADA model with and without adjustments.

5. Importance of maximum range versus maximum endurance

As described earlier, the fuel flow per unit time should be minimized when time constraints exist. Figure 8b shows that fuel flow per unit time has less variation with altitude than fuel flow per unit distance does, for nominal cruise airspeed profiles. However, the aircraft's nominal holding speed, rather than its nominal cruise speed, is a more appropriate representation of the actual speed flown when time, rather than distance, is constrained. Nominal cruise speeds are usually slightly faster than the aircraft's maximum range speed (i.e., fuel flow per unit distance could be decreased by flying slower). Therefore, they are also significantly faster than the aircraft's maximum endurance speed (i.e., fuel flow per unit time could be decreased by flying even slower). Ultimately, the efficiency of absorbing delay at higher altitudes, when it must be absorbed in the form of level segments, is bounded by the variation of the aircraft's maximum endurance speed with altitude. The recommended holding speeds for commercial jet aircraft are slightly faster than their maximum endurance speeds in order to provide adequate maneuver margin, ensure passenger comfort, and minimize throttle activity. The recommended holding speeds for commercial jet aircraft, at typical top-of-descent weights, are between 210 KIAS and 250 KIAS. The recommended speeds are strongly affected by aircraft weight, but vary little with altitude. Figure 9 shows the estimated fuel flow predicted by BADA for a Boeing 737-800 (59,800 kg, ISA, FLAPS UP) for a series of delay absorbing airspeed profiles. Those speed profiles range from the nominal cruise speed profile (fastest) to the recommended holding speed profile (slowest). These curves suggest three simple principles in the choice of speed and altitude:

1. The fuel flow rate per unit time is reduced slightly by flying higher.
2. The fuel flow rate per unit time is reduced significantly by flying slower.
3. The fuel flow rate for the nominal cruise speed profile is considerably greater than the fuel flow rate for the recommended holding speed profile.

The underlying insight is that the airspeed profile becomes more important to fuel efficiency than the altitude profile when delay must be absorbed using level segments.

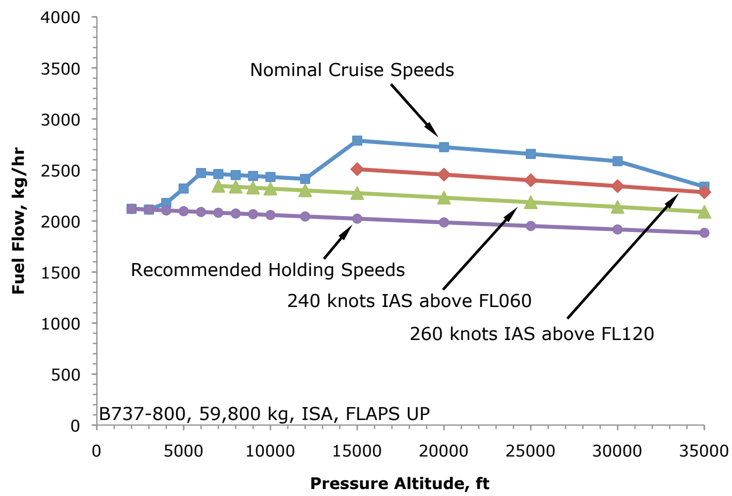

Figure 9. Estimates of fuel flow per unit time for BADA model with adjustments for various delay absorbing speed profiles.

V. Results

Potential fuel savings, for uncongested (i.e., distance-constrained) and congested (i.e., time-constrained) operations, for every flight in every traffic sample were calculated. Estimates, with and without fuel flow corrections, were calculated to quantify the effects of the fuel model adjustments. In addition, potential fuel savings were estimated for various initiation and termination altitudes in order to understand the benefits of different continuous descent scenarios. In order to better understand the nature of today's terminal operations, the vertical inefficiencies for each flight, in terms of the length of level flight segments flown below 17,000 feet MSL and the associated fuel savings, was calculated. This section will discuss those results in detail.

1. Variability of potential fuel savings between flights and between days

Prior studies have reported the benefits of continuous descents as average fuel savings per flight and total fuel savings. However, these simple statistical measures do not adequately describe the benefits of continuous descents for the typical flight.

The variability between flights of the potential fuel savings was calculated in order to understand how the diversity of the traffic sample affected the results. Analysis of the potential fuel savings per flight, for the distance-constrained scenario, shows that its statistical distributions for each airport exhibit long tails. In general, the mean fuel savings per flight is 20-80% greater than the median fuel savings for the same traffic sample. One cause of the disparity between the mean and median values is that fuel flow rates vary by aircraft type. For example, fuel flow rates for aircraft of the super-heavy and heavy weight-classes are 2-4 times larger than those of aircraft of the large weight-class, and more than eight times larger than those of regional and business jets. Thus, whereas aircraft of the super-heavy and heavy weight-class represent approximately 8% of all operations, they account for approximately 22% of the potential fuel savings. Limited analyses suggest that a sample reduced to a single weight-class on the same arrival route to the same runway exhibits reduced disparity between the mean and median per flight values. In order to represent the benefit of a continuous descent to the typical flight, the median value, rather than its mean value, will be reported. Ultimately, the use of either statistic will depend upon whether the purpose is to achieve the greatest total fuel savings, the most equitable distribution of fuel savings across all flights, or somewhere in between.

The variability between traffic samples of the potential fuel savings for each airport was also measured in order to understand how the number of daily traffic samples affected the results. The median fuel savings per flight of an airport's daily traffic samples was compared to the median fuel savings of all of its traffic samples combined. The standard deviation between traffic samples of the median fuel savings per flight was approximately 20%. The median fuel savings for an airport's daily traffic samples was generally within ±40% of the median fuel savings for all of its traffic samples combined. Thus, the potential fuel savings of the airport's worst day appeared twice as large as that of its best day. Most airports exhibited similar variability between days. However, the NCT airports (OAK, SFO and SJC) had almost twice the variability between days as the other airports had. It is not known whether this variability is caused by different airspace and airport configurations, or simply by the day-to-day differences of operations for the same airspace and airport configuration.

2. Comparisons of potential fuel savings for different fuel flow model adjustments

Comparisons of the median fuel savings for the different fuel model adjustments were made for each airport. The median fuel savings per flight was calculated for the distance-constrained trajectories for all of the traffic samples using the full continuous descent scenario that eliminated all level segments below cruise.

Figure 10 shows a comparison of the median fuel savings per flight for the different fuel model adjustments. The blue portions of the columns show the potential fuel savings predicted by the original BADA fuel flow model with no adjustments. The red portions of the columns show the effect of the cruise fuel flow rate correction. The green portions show the effect of reducing the nominal cruise speed in terminal airspace. And, finally, the purple portions show the effect of modeling aircraft configurations using flaps at low altitudes. In general, the median fuel savings per flight is doubled by application of the model adjustments described in Sec. IV-C. The correction of the cruise fuel flow rate accounted for 50-75% of the increase. The reduction of the nominal cruise speed increased the median fuel savings per flight by a modest 5-15% depending on the airport. Finally, the modeling of the FLAPS 1 and FLAPS 5 aircraft configurations accounted for a further increase of 15-25%.

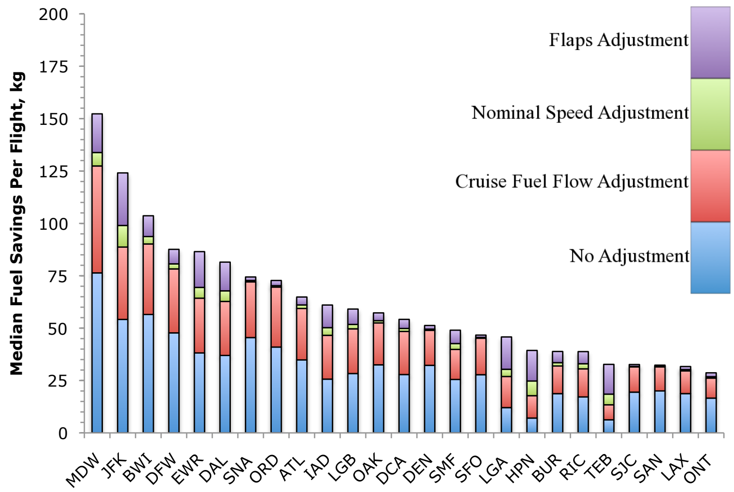

Figure 10. Effect of fuel flow model adjustments on median fuel savings per flight by airport.

The original BADA fuel flow model does not predict the effects of level flight at lower altitudes with sufficient accuracy. The sensitivity of the potential fuel savings to the fuel flow model is shown to be significant; whereas the increase in fuel flow rate per unit time at all altitudes did not exceed 50%, the potential fuel savings per flight was doubled for most airports. By contrast, the relative ranking of the airports by their potential fuel savings was not significantly affected by the adjustments.

3. Comparison of potential fuel savings for full and partial continuous descents

Comparisons of the median fuel savings for various full and partial continuous descent scenarios were made for each airport. The median fuel savings per flight was calculated for the distance-constrained trajectories of all traffic samples using the adjusted fuel flow model. Figure 11 shows the results for the full continuous descent scenarios that replaced level segments below the different key altitudes with level segments in cruise. The colored portions of the columns show the amount of savings associated with level segments in different altitude ranges. The color-coding scheme was as follows: blue for terminal airspace, red for transition airspace, green for en route low sectors and purple for en route high sectors. The terminal, transition, low en route and high en route level segments account for 62, 20, 8 and 10% of the potential fuel savings, respectively. The results for DEN match the findings of Ref. 13.

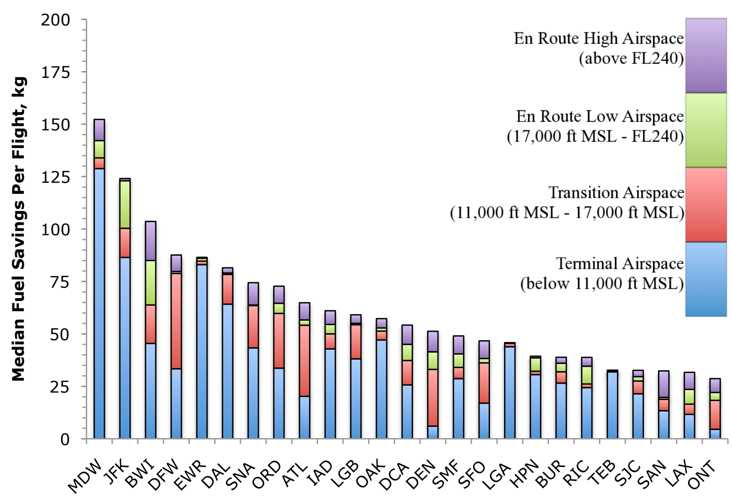

Figure 11. Effect of continuous descent scenario on median fuel savings per flight by airport.

The distributions shown for the different airports exhibit markedly different behaviors. For example, level segments in terminal airspace account for more than 80% of the potential fuel savings at DAL, MDW, and OAK.

Much of this level distance is a result of delay absorption in the TRACON and procedural separation from aircraft to the nearby dominant airport (i.e., DFW, ORD, and SFO). Elimination of these level segments would require improved scheduling, as well as airspace redesign. Meanwhile, ATL, DFW, ORD and SFO have similar shares of potential fuel savings for level segments in both terminal and transition airspace. At these airports, long level segments upon entry into the TRACON account for more than 33% of the potential fuel savings. Elimination of these level segments might be possible by changing existing arrival procedures to raise the arrival meter fix's crossing altitude.

The results for the partial continuous descent scenarios, whose level segments were replaced with level segments at the key altitudes instead of the cruise altitude, showed approximately half of the savings as the full continuous descent scenarios. For example, the scenario replacing all level segments in terminal airspace with level segments at 11,000 feet MSL instead of cruise altitude showed 55% less potential fuel savings. While this benefit reduction is substantial, it does indicate that partial continuous descent scenarios (i.e., ones that initiate the continuous descent from higher intermediate altitudes instead of from cruise altitude) should be investigated as an alternative if a costly airspace redesign or infrastructure improvements must be avoided.

4. Comparison of potential fuel savings for distance-constrained and time-constrained trajectories

Level segments after descent from cruise altitude are associated with a time penalty, as well as a fuel penalty. Elimination of the level segments at intermediate altitudes decreases the flight time from top-of-descent to landing, because flights are able to maintain faster speeds at the higher altitudes. The median time savings per flight was calculated for the distance-constrained trajectories of all traffic samples using the adjusted fuel flow model. The result was a median flight time savings of approximately two minutes. The 25th percentile and 75th percentile values were approximately one minute and three minutes, respectively.

Figure 12 shows the effect of the time-to-fly constraint on the median fuel savings per flight for each airport. The blue columns show the fuel savings associated with the elimination of all level segments below cruise. The red columns show the fuel savings of the same scenario when the time constraint is enforced in the manner described above. The estimated fuel savings per flight for the time-constrained trajectories were 70-85% less than the estimated fuel savings per flight for the distance-constrained trajectories.

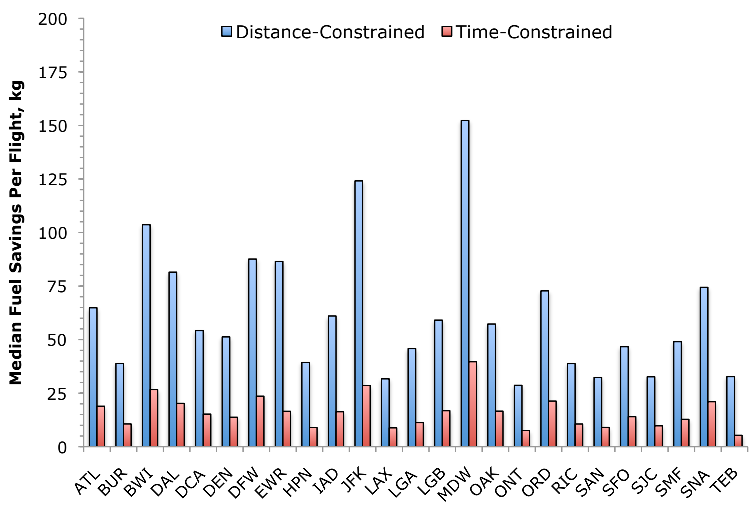

Figure 12. Comparison of median fuel savings per flight for distance-constrained (uncongested) and time-constrained (congested) trajectories

As stated previously, these results represent a worst-case scenario, because they do not allow the delay to be absorbed using descent speed reductions. As described in Sec. IV-C, air traffic controllers would reduce cruise and descent speed to absorb some fraction of the required delay. Slower descent speeds would result in different, yet still near idle, descent profiles, and would not result in a reduction of the potential fuel savings. However, as shown in Fig. 9, the fuel flow rate per unit time exhibits only a small variation with altitude for level flight at similar speeds and aircraft configurations. As a result, any necessary constant-time level segments will be only slightly more fuel efficient at cruise altitude than at lower altitudes. Therefore, the effect of the time constraint is significant. It shows that the potential fuel savings must necessarily be reduced once delay needs to be absorbed in the form of level segments. These results will appear counterintuitive if only the variation with altitude of the fuel flow rate per unit distance is considered because the former variation is substantial. The reduction of the fuel savings would be most severe when the delay required to be absorbed is the greatest.

5. Characterization of vertical inefficiencies across multiple airports

The rest of the paper will focus on analyses of the level segments in transition and terminal airspace (i.e., below 17,000 feet MSL). These level segments accounted for approximately 80% of the estimated fuel savings. The results, when reported for all airports, have been adjusted to account for the different numbers of traffic samples analyzed for each airport.

Figure 13 shows the median level distance below 17,000 feet MSL per flight for each airport. The bars indicate the 25th percentile (shorter) and 75th percentile (longer) distances. Typically, flights flew level below 17,000 feet MSL for 20-30 nmi. The level distance was less than 25 nmi for 52% of the flights, but it was more than 55 nmi for 8% of the flights. Again, the probability distribution of the level distance per flight exhibits a long tail, and varies significantly by airport. The median level distance was the greatest for the C90 and N90 airports; it was the least for the NCT and SCT airports. For example, only 16% of SFO flights had level distances below 17,000 feet MSL greater than 30 nmi, whereas 58% of ORD flights did. These results show, in particular, that extrapolation of seemingly comparable airports on the basis of simple characteristics (e.g., number of operations, perceived complexity, etc.) does not capture the uniqueness of each airport's operations. A similarity between the airports, however, was that the difference between the 25th percentile and 75th percentile values was nearly uniform.

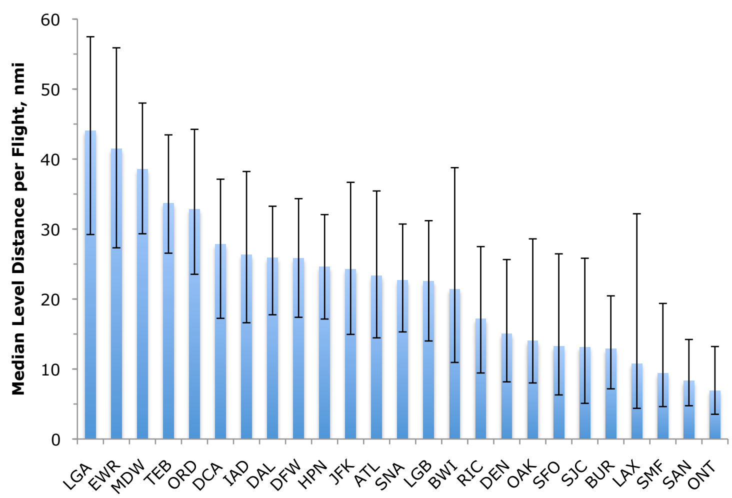

Figure 13. Median level distance flown below 17,000 feet MSL per flight by airport.

Additional analyses showed that the median level distance below 17,000 feet MSL increased by approximately 5-15 nmi between low traffic demand (fewer than 16 arrivals per runway per hour) and high traffic demand (more than 32 arrivals per runway per hour). For some airports, the level distance variation between moderate and high traffic demand did not exhibit a strong trend. In particular, most of the level distance increase occurred between low and moderate traffic demand for ATL, DFW, EWR, MDW, and ORD. These results suggest that some controllers use much of their available path stretch options even for moderate traffic demand. Thus, extrapolation of potential fuel savings from moderate traffic demand (for which field trials of continuous descent procedures do exist) to high traffic demand by scaling level distance would likely overestimate the potential benefit.

The adjusted fuel flow model was used to estimate the potential fuel savings associated with the level segments characterized in the preceding paragraphs. Figure 14 shows the median fuel savings per flight for each airport. Again, the bars indicate the 25th percentile (lower) and 75th percentile (higher) savings. For comparison with Fig. 13, the airports are ordered by median level distance per flight. Across airports, there is little correlation between the level distance per flight and the fuel savings per flight. The relationship of potential fuel savings to level distance is a complex combination of the flight's aircraft type and the altitude of the level segment in relation to the flight's cruise altitude. Airports with large fractions of heavy aircraft, like JFK, will naturally have higher fuel savings for shorter level distances. Meanwhile, airports with smaller fractions of heavy aircraft, like LGA, will have more moderate fuel savings for longer level distances.

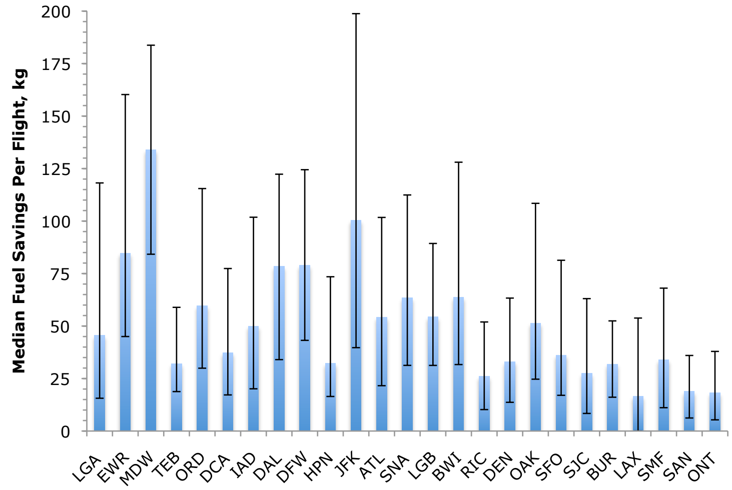

Figure 14. Median fuel savings associated with level segments flown below 17,000 feet MSL per flight by airport.

In general, the estimated fuel savings per flight is quite modest. The estimated fuel savings per flight was less than 100 kg for more than 87% of the flights. In fact, the estimated fuel savings per flight was less than 25 kg for more than 45% of flights. The 13% of the flights with the greatest savings accounted for 37% of the estimated fuel savings.

The potential fuel savings for the N90 airports (EWR, JFK and LGA) and PCT airports (BWI, DCA and IAD) were likely underestimated by 10-20% with respect to the other airports. The combination of en route flow constraints and close proximity to ARTCC boundaries caused the apparent cruise altitude of many flights to these airports to appear low (i.e., below FL240). Coordination of continuous descents through adjacent facilities and consideration of each flight's desired cruise altitude could increase the potential savings. However, determination of a more precise estimate would need to consider the flight's stage length, since these airports also have many very-short-haul commuter flights.

VI. Discussion

As described earlier, the variability between flights and between days of the potential fuel savings, as well as the small sample sizes of previous studies, make detailed cross-study comparisons difficult. Comparisons of simple statistical measures suggest that the results of this study are consistent with the estimated fuel savings reported in previous work. The previous studies that analyzed traffic samples for individual airports (Ref. 12 through 15) estimated the average fuel savings per flight to be between 54 and 157 kg for full continuous descents from cruise altitude. The average fuel savings per flight estimated by this study was 98 kg (see Fig. 14). Ref. 6 predicted that a Boeing 777-200 aircraft would save approximately 110, 160, and 1462 kg of fuel during periods of low, moderate and high congestion, respectively. The average fuel savings per Boeing 777-200 flight estimated by this study, for all traffic demand levels, was 263 kg.

For all flights in the traffic samples studied, using the distance-constrained trajectories, the potential fuel savings associated with the level segments below 17,000 feet MSL was estimated to be 3% of their total fuel consumption. Other studies have found similar benefits as a proportion of the total fuel consumption.[15] These studies have also scaled the fuel savings per flight to the entire NAS. Invariably, the potential NAS-wide yearly fuel savings for routine continuous descent operations is large. However, a discussion of the impact of these aggregate benefits, in the context of their effect on the flying public and their associated implementation costs, is important.

First, the expected fuel savings per flight is not likely to reduce the cost of flying for the passengers. One gallon of jet fuel weighs approximately 3.1 kg (6.84 lb), and at recent fuel prices, costs approximately $2.20 USD. Thus, a 150 kg (331 lb) fuel savings for a typical single-aisle, narrow-body jet aircraft would be approximately $1 per passenger. Furthermore, the fuel savings of the level segments below 17,000 feet MSL account for less than 1% of the total fuel consumption for the 45% of flights whose savings was predicted to be less than 25 kg. Second, the infrastructure improvements needed to conduct continuous descents during periods of high traffic demand are significant. All forms of continuous descents require the development of new arrival procedures, and likely the redesign of existing airspace. Furthermore, they have an implicit assumption that improved scheduling capabilities on the ground limit the amount of unexpected delay absorption needed past the initial top-of-descent. The costs of these improvements will be significant in order to realize the full benefits of continuous descents. Finally, redesigned procedures and better scheduling, by themselves, will eliminate some of the level distance flown after initial top-of-descent, and thus achieve some of the improvements often associated specifically with full continuous descents.

VII. Conclusions

The potential fuel savings of continuous descents was estimated for more than 480,000 arrival flights for 25 major airports at eight busy TRACONs in the NAS. Between thirty and sixty days were analyzed for each airport. Baseline trajectories were constructed from recorded flight plan and track data, and their level segments were identified using a set of pattern-matching heuristics. One set of continuous descent trajectories was calculated by moving level segments to higher altitudes while enforcing a constant distance-to-fly constraint to simulate uncongested operations. Another set of continuous descent trajectories was calculated by moving level segments to higher altitudes while enforcing a constant time-to-fly constraint to simulate congested operations. BADA was used to estimate the fuel and flight time savings associated with changes to the vertical profiles of the baseline trajectories. The results of these analyses were reported for each airport in terms of probability distributions and the sensitivity of these distributions for different aircraft weight-classes, traffic demand levels, and continuous descent scenarios. Four broad conclusions are supported by these analyses.

First, the estimated fuel savings are found to be sensitive to the size (e.g., number of days, flights, etc.) and diversity (e.g., proportion of different weight-classes, arrival routes, etc.) of the traffic samples analyzed. The potential fuel savings of an airport's worst day appeared twice as large as that of its best day. The variability between flights of the potential fuel savings illustrates that simple statistical measures (i.e., mean and total fuel savings) do not adequately describe the fuel savings of continuous descents for the typical flight. In particular, aircraft of the super-heavy and heavy weight-classes were approximately 8% of the total operations, but they were approximately 22% of the predicted fuel savings.

Second, the estimated fuel savings are significantly affected by the choice of fuel flow model and continuous descent scenario. Relatively small adjustments to the BADA fuel flow predictions had substantial impact on the predicted fuel savings. Whereas the adjustments increased the estimated fuel flow rate less than 50%, the predicted fuel savings, in response to these adjustments, was doubled for most airports. The adjustment associated with the use of flaps indicates that detailed consideration of the aircraft configuration is needed to accurately estimate potential benefits. It also indicates that speed profiles that ensure FLAPS UP operations, regardless of vertical profile, will have considerable fuel savings. As expected, the potential fuel savings is greatest for the elimination of level segments at the lowest altitudes. Level segments in the terminal, transition, low en route and high en route airspace account for 62, 20, 8 and 10% of the potential fuel savings, respectively. However, different airports exhibit markedly different distributions. Thus, extrapolation of the estimated fuel savings of one airport to other seemingly comparable airports, on the basis of simple characteristics (e.g., number of operations, perceived complexity, etc.), will not capture the uniqueness of each airport's operations, or hence its potential fuel savings.

Third, the typical flight's estimated fuel savings associated with eliminating level segments below 17,000 feet MSL was modest. The duration of these level segments was less than 25 nmi for 52% of the flights. Also, the level distance showed considerable variation between airports; only 16% of SFO flights had level distances greater than 30 nmi, while 58% of ORD flights did. Elimination of these level segments accounted for approximately 80% of the fuel savings estimated for full continuous descents from cruise. Furthermore, the estimated fuel savings was not distributed evenly among flights. It was less than 25 kg for more than 45% of flights, and less than 100 kg for more than 87% of the flights. As a result, 13% of the flights accounted for 37% of the total fuel savings.

Lastly, the estimated fuel savings were not found to scale with traffic demand, and they showed diminished returns when delay needed to be absorbed in the form of level segments. The median level distance below 17,000 feet MSL increased by approximately 10 nmi between low and high traffic demand conditions. However, it did not show a strong trend between moderate and high traffic demand conditions for some airports. Enforcement of the time constraint on the continuous descent trajectories reduced estimated fuel savings by up to 70-85%. The reduction of potential benefits would be most severe when the delay required to be absorbed was the greatest.

Finally, continuous descents have demonstrable fuel savings, especially when executed during periods of low to moderate traffic demand. Many studies have consistently indicated that the potential fuel savings per flight is between 50 and 150 kg. However, the magnitude of these benefits, and their importance, depends upon the context in which they are described. For example, the fuel savings associated with continuous descents, as a fraction of total fuel consumption, was estimated to be approximately 3%. This fuel savings equates to less than $1 USD per passenger for a typical single-aisle, narrow-body jet aircraft. Ultimately, the NAS improvements needed to conduct continuous descents during periods of high traffic demand will incur substantial costs to develop advanced ground automation, create new arrival procedures, and redesign existing airspace structures. Therefore, the prioritization of these improvements must rely upon an analysis of a sufficient sample of operations that is representative of many days, aircraft types, and traffic demand levels. This prioritization must use detailed models of aircraft configuration, airspeed profiles and fuel flow, and must consider alternate continuous descent scenarios that do not eliminate all level segments.

Acknowledgements

The authors would like to thank Michael Haney of San Jose State University, and Christopher Wu of Archbishop Mitty High School for supporting the software development and data analysis efforts of this research. The authors are also grateful for the valuable feedback, insight, and feedback provided by Doug Isaacson, Harry Swenson and Steve Green of NASA; and Jack Dwyer of the Boeing Company.

References

[1] Joint Planning and Development Office, "Concept of Operations for the Next Generation Air Transportation System," Version 2.0, June 2007.

[2] Isaacson, D. R., Swenson, H. N., Robinson III, J. E., "A Concept for Robust, High Density Terminal Air Traffic Operations," 10th AIAA Aviation Technology, Integration, and Operations Conference (ATIO), Fort Worth, TX, Sep. 13-15, 2010, To Be Published.

[3] White, W., and Clarke, J.-P., B., "Details and Status of CDA Procedures at Los Angeles International Airport (LAX)," Presented at *CDA Workshop No. 3*, Georgia Institute of Technology, Atlanta, GA, September 6-7, 2006.

[4] Staigle, T., and Nagle, G., "Details and Status of CDA Procedures for Early Morning Arrivals at Hartsfield-Jackson Atlanta International Airport (ATL)," Presented at *CDA Workshop No. 3*, Georgia Institute of Technology, Atlanta, GA, September 6-7, 2006.

[5] Clarke, J.-P. B., Ho, N. T., Ren, L., Brown, J. A., Elmer, K. R., Tong, K.-O., and Wat, J. K., "Continuous Descent Approach: Design and Flight Test for Louisville International Airport," *Journal of Aircraft*, Vol. 41, No. 5, 2004, pp. 1054-1066.

[6] Coppenbarger, R. A., Mead, R. W., Sweet, D. N., "Field Evaluation of the Tailored Arrivals Concept for Datalink-Enabled Continuous Descent Approach," *Journal of Aircraft*, Vol. 46, No. 4, 2009, pp. 1200-1209.

[7] Tong, K.-O., Boyle, D. A., and Warren, A. W., "Development of Continuous Descent Arrival (CDA) Procedures for Dual-Runway Operations at Houston Intercontinental," *6th AIAA Aviation Technology, Integration, and Operations Conference (ATIO)*,Wichita, KS, Sept. 25-27, 2006, AIAA-2006-7750.

[8] Wat, J., Follet, J., Mead, R., and Brown, J., "In Service Demonstration of Advanced Arrival Techniques at Schipol Airport," *6th AIAA Aviation Technology, Integration, and Operations Conference (ATIO)*,Wichita, KS, Sep. 25-27, 2006, AIAA-2006-7753.

[9] Reynolds, T. G., "Development of Flight Inefficiency Metrics for Environmental Performance Assessment of ATM", *8th USA/Europe Seminar on Air Traffic Management Research and Development (ATM2009)*, Napa, CA, June 29-July 2, 2009, Paper No. 122.

[10] "User Manual for the Base of Aircraft Data (BADA)", Revision 3.7, EEC Technical/Scientific Report No. 2009-003, Eurocontrol Experimental Centre, March 2009.

[11] Sprong, K. R., Klein, K. A., and Shiotsuki, C., "Analysis of AIRE Continuous Descent Arrival Operations at Atlanta and Miami," *27th Digital Avionics Systems Conference (DASC)*, St. Paul, MN, Oct. 26-30, 2008, pp. 3.A.5-1-3.A.5-13.

[12] Dinges, E., "Determining the Environmental Benefits of Implementing Continuous Descent Approach Procedures," *7th USA/Europe Seminar on Air Traffic Management Research and Development (ATM2009)*, Barcelona, Spain, July 2-5, 2007, Paper No. 594.

[13] Shresta, S., Neskovic, D., and Williams, S. S., "Analysis of Continuous Descent Benefits and Impacts During Daytime Operations," *8th USA/Europe Seminar on Air Traffic Management Research and Development (ATM2009)*, Napa, CA, June 29-July 2, 2009, Paper No. 132.

[14] Melby, P., Mayer, R. H., "Benefit Potential of Continuous Climb and Descent Operations," *The 26th Congress of International Council of Aeronautical Sciences (ICAS)*, Anchorage, AK, Sep. 14-19, 2008, AIAA-2008-8920.

[15] Alcabin, M., Schwab, R. W., Soncrant, C., Tong, K.-O., and Cheng, S. S., "Measuring Vertical Flight Path Efficiency in the National Airspace System," *9th AIAA Aviation Technology, Integration, and Operations Conference (ATIO)*, Hilton Head, SC, Sep. 21-23, 2009, AIAA-2009-6959.

[16] De Los Santos, V., "Research Foundations Main," *NASA Ames Aviation Systems Division*, NASA Ames Research Center, 12 Dec. 2008, Web, 4 Feb. 2010, <http://www.aviationsystemsdivision.arc.nasa.gov/research/foundations/index.shtml>.

[17] "Airport Development Solution Set Smart Sheet," Federal Aviation Administration, 3 July 2008, Web, 25 Aug 2010, <http://www.faa.gov/about/office_org/headquarters_offices/ato/publications/nextgenplan/0608/solution_sets/oep_35/>.

[18] Air Traffic Control FAA Order JO 7110.65S, Federal Aviation Administration, Washington, DC. Feb. 2008. (Published online at http://www.faa.gov).

[19] "Airport Surface Detection Equipment, Model X (ASDE-X)," Federal Aviation Administration, 26 June 2009, Web, 25 Aug 2010, <http://www.faa.gov/air_traffic/technology/asde-x/>.

[20] "Automated Radar Terminal System," Federal Aviation Administration, 29 June 2009, Web, 25 Aug 2010, <http://www.faa.gov/air_traffic/technology/tamr/arts/>.

[21] Nuic, A., Poinsot, C., Iagaru, M.-G., Gallo, E., Navarro, F. A., and Querejeta, C., "Advanced Aircraft Performance Modeling for ATM Enhancements to the BADA Model," *24th Digital Avionics Systems Conference (DASC)*, Washington, DC, Oct. 30-Nov. 3, 2009, pp. 2.B.4-1-2.B.4-12.

[22] Senzig, D. A., Fleming, G. G., and Iovinelli, R. J., "Modeling of Terminal-Area Airplane Fuel Consumption", *Journal of Aircraft*, Vol. 46. No. 4, 2009, pp. 1089-1093.

[23] Roach, K., and Robinson III, J. E., "A Terminal Area Analysis of Departure Level-Offs at Dallas/Fort Worth International Airport," 10th AIAA Aviation Technology, Integration, and Operations Conference (ATIO), Fort Worth, TX, Sep. 13-15, 2010, To Be Published.

[24] De Los Santos, V., "CVSRF Main," *NASA Ames Aviation Systems Division*, NASA Ames Research Center, 12 Dec. 2008, Web, 25 Aug. 2010, <http://www.aviationsystemsdivision.arc.nasa.gov/facilities/cvsrf/index.shtml>.

[25] Continental Airlines Flight Manual for B777-224, Revision #17, 9 Aug. 2002.