7th AIAA Aviation Technology, Integration and Operations Conference (ATIO)
18 - 20 September 2007, Belfast, Northern Ireland

AIAA 2007-7778

Field Evaluation of the Tailored Arrivals Concept for Datalink-Enabled Continuous Descent Approach

Richard A. Coppenbarger*
NASA Ames Research Center, Moffett Field, CA, 94035

Rob W. Mead†
The Boeing Company, Seattle, WA, 98124

and

Douglas N. Sweet‡
Sensis Corporation, Campbell, CA, 95008

Allowing aircraft to descend uninterrupted at low engine power, continuous descent operations promise to maximize fuel efficiency while minimizing environmental impact. Tailored Arrivals is a concept for enabling continuous descents under constrained airspace conditions by integrating advanced air and ground automation through digital datalink. Operational trials were completed in January 2007 involving trans-pacific flights into San Francisco during early morning hours. Leveraging newly deployed FAA automation in the Oceanic environment, trajectory-based clearances were transmitted by datalink to Boeing 777 aircraft equipped with Future Air Navigation System avionics. NASA's prototype, ground-based automation for high-density arrival management tailored trajectory clearances to accommodate artificially imposed metering constraints. Upon sharing wind and descent-speed-intent data, ground-based and airborne automation were found to predict meter-fix arrival times to within a mean accuracy of 3 seconds over a 25-minute prediction horizon. Corresponding mean altitude and along-track prediction errors of ground-based automation were -500 ft and -1.3 nmi, respectively, in comparison to surveillance truth. A benefits analysis suggests Boeing 777 fuel savings of between 200 and 3,000 lbs per flight – depending highly upon baseline traffic conditions – together with a corresponding reduction in CO_2 emissions of between 700 and 10,000 lbs per flight.

I. Introduction

ALTHOUGH airborne capabilities such as Area Navigation (RNAV) and Required Navigation Performance (RNP), together with optimized guidance and control through the Flight Management System (FMS), offer substantial improvements to the efficiency of flight operations, their benefits frequently go unrealized in today's arrival airspace domain. To the frustration of the airspace user, efficiency gained en route using state-of-the-art airborne automation is often squandered during the final stages of flight as the airplane transitions for landing. Here, Air-Traffic Control (ATC) actions often require the airplane to execute sub-optimal, tactical maneuvers that involve frequent temporary-altitude assignments, speed adjustments and lateral vectoring to accommodate tightly coupled separation and traffic-flow-management constraints. These actions, though designed to manage controller workload and ensure safety in heavy traffic conditions, prevent aircraft from executing an uninterrupted Continuous Descent Approach (CDA) to the runway using low engine power for maximum fuel efficiency and minimum environmental impact. As a result, FMS guidance trajectories, which are optimized over time horizons that transcend ATC boundaries, are seldom executed to completion in today's arrival airspace. Accommodating efficient, trajectory-

* Aerospace Engineer, Automation Concepts Research Branch, Mail Stop 210-10, AIAA Senior Member
† Lead Engineer –Air/ground Comm and Tailored Arrivals, Boeing Phantom Works Advanced ATM, P.O. Box 3707 MC 8R-79
‡ Director, Sensis Seagull Technology Center, 1700 Dell Ave., AIAA Member

based arrival operations under all traffic conditions is a key objective of the Next Generation Air Transportation System (NextGen).[1]

The benefit and feasibility of CDA operations has been a subject of considerable study.[2,3] These studies have occasionally involved flight trials, some evolving into limited daily use operations into select airports around the world. Examples include the recent Advanced Arrivals trials in the Netherlands,[4] CDA trials at Louisville,[5,6] flight trials in the U.K.[7] and Sweden, and the initial Tailored Arrivals trials conducted in Australia.[8] Although these studies have involved a myriad of techniques for enabling CDA benefits, their focus has largely been on developing static arrival procedures for near-term deployment. As a result, the application of these CDA initiatives requires either periods of low traffic density or specialized protocols that limit throughput.

A key feature of the flight trials described in this paper, which distinguishes them from related activities in the U.S. and abroad, is the inclusion of ground-based automation capable of generating dynamic CDA trajectory solutions in the presence of complex airspace constraints. The ability to tailor arrival solutions to accommodate individual aircraft performance, atmospheric conditions, and operational restrictions is critical for enabling CDAs in congested airspace environments where potential benefits are greatest. Towards this objective, NASA's En Route Descent Advisor (EDA) was incorporated into the trials to compute advisories for meeting arrival-time constraints artificially imposed at the TRACON boundary, designed to emulate operations under heavy traffic conditions. EDA is a research component of the Center-TRACON Automation System (CTAS) that works in conjunction with the CTAS Traffic Management Advisor (TMA) to provide controllers with combinations of speed, altitude and path-stretching advisories that satisfy time-based metering constraints while avoiding separation conflicts.[9-11]

The automation systems, clearance composition, and human procedures used in carrying out the EDA-supported OTA trials are first described, followed by results describing potential fuel and emissions benefits. Due to the importance of accurate ground-based and airborne trajectory predictions in planning and executing Tailored Arrivals, results are then presented that compare EDA and FMS trajectory-prediction performance. Additional results obtained from these trials pertaining to human procedures, noise exposure, and wind-data comparisons will be documented in detail at a later date.

II. Approach

In collaboration with the FAA and United Airlines, The Oceanic Tailored Arrivals (OTA) trials were conducted over 40 days with a single United Airlines Boeing 777 flight (UAL76) in commercial service between Honolulu and San Francisco (SFO). The flight was operated along the Central-East-Pacific (CEP) oceanic route structure, comprised of a series of fixed, parallel tracks from the Hawaiian Islands to the California coast.[12] In addition to its avionics equipage, UAL76 was chosen largely because of its early morning (5:30 AM local) arrival time at SFO, which avoided congested airspace conditions, thereby minimizing the likelihood of interference with other traffic.

A. System Components

Figure 1 illustrates the key systems employed in formulating, communicating, and executing the trajectory-based arrival clearances used to support the OTA field trials (the precise content of the data communications between the various system elements is described later in the paper). Clearance delivery was initiated approximately 700 nmi from landing in oceanic airspace controlled by the Oakland Air-Route Traffic Control Center (ARTCC), referred to as ZOA. In oceanic airspace, controllers at ZOA rely upon the FAA's recently deployed Advanced Technologies and Oceanic Procedures (ATOP) system for integrated CNS functions. A prominent capability of ATOP is its ability to support two-way digital messaging between ATC and the flight deck through Controller-Pilot Data-link Communications (CPDLC). The OTA route clearance – comprised of lateral waypoints together with speed and altitude restrictions – was relayed to the flight deck using a standard CPDLC message format supported by ATOP. The format allowed the OTA instructions, upon pilot approval, to be directly loaded into the B777 FMS through its Future Aircraft Navigation System (FANS) avionics interface. FANS avionics integrate FMS functionality with CPDLC and contract-based Automatic Dependent Surveillance (ADS-C) services in oceanic airspace.

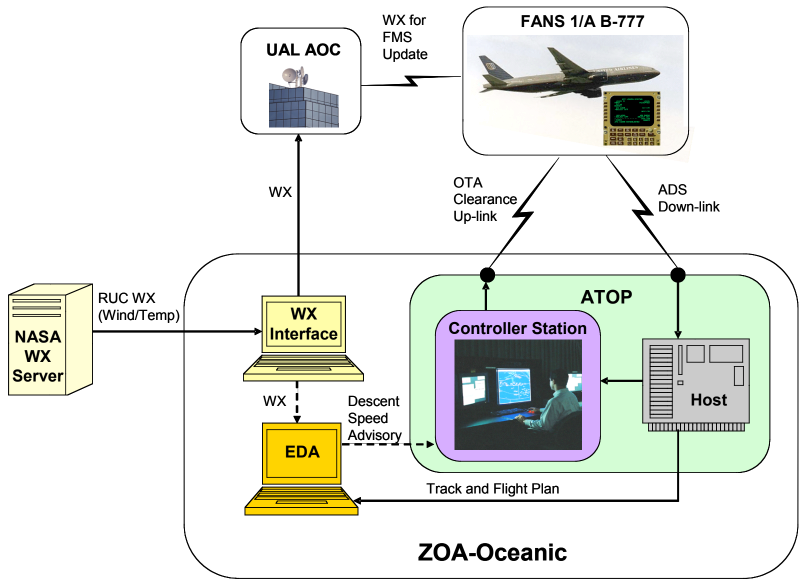

Figure 1. OTA System Components

Once loaded, the OTA route clearance provided sufficient information for the FANS FMS to compute a 4-D reference trajectory from the aircraft's current position to the runway. This trajectory was then used by the FMS as the basis for its Lateral Navigation (LNAV) and Vertical Navigation (VNAV) guidance functions, which provided inputs to the automatic flight control system for determining appropriate aileron, rudder, elevator and throttle inputs. Once activated, the OTA route clearance allowed the flight to progress with no additional pilot inputs required prior to configuring the airplane for landing. The nominal guidance law used by the FMS VNAV function was set to PATH mode, which directed the autopilot to null any positional error between the airplane's current altitude and the FMS reference trajectory's vertical profile. In the event that the airspeed required to control path differed from that suggested by the reference trajectory by more than ±10 kt, the VNAV guidance laws would switch from PATH to SPEED control mode.

To provide a dynamic element to the OTA trajectory-based clearance, a prototype version of NASA's EDA decision-support tool was incorporated into the field trials. EDA was used to compute the maneuver solution needed to target a meter-fix crossing time constraint imposed at the Terminal Radar Approach Control (TRACON) boundary. Originally designed to assist the domestic en route sector controller in developing conflict-free arrival metering solutions under capacity constrained conditions, EDA was adapted to ZOA airspace and interfaced with ATOP to receive oceanic surveillance and flight-plan data. This oceanic surveillance data was derived from the airplane's satellite-based positioning system, and relayed to ATOP via ADS-C at the maximum-available rate of once every 2 minutes, as specified in the ADS-C contract configured by ZOA. In addition to surveillance and flight plan inputs, EDA required atmospheric data to compute the long-range trajectory predictions needed for its advisories. These input data were derived from the National Oceanic and Atmospheric Administration's Rapid Update Cycle (RUC) model, which provided 2-hour forecasts (updated each hour) of wind speed, wind direction, temperature and pressure, organized in a Lambert-Conformal 3-D grid with a lateral resolution of 40 km.

For the purpose of the OTA trials, EDA computed only descent Calibrated Airspeed (CAS) advisories for targeting a Scheduled Time of Arrival (STA) at the TRACON meter fix (waypoint BRINY). EDA cruise-speed and path-stretch advisory capabilities, along with automated conflict resolution functions, were purposely suppressed to avoid complexity and unnecessary risk in these initial trials. For optimal TRACON throughput under congested conditions, EDA derives its target STAs from TMA. During the time of these trials, however, no capacity constraints existed, so meter-fix STAs were artificially set to within ±4 minutes of the airplane's original Estimated Time of Arrival (ETA) at BRINY. This original ETA was computed from an EDA trajectory prediction using a nominal assumption of company-preferred descent speed (280 kt).[11] An example of EDA's graphical user interface, showing a descent CAS advisory for UAL76 of 267 kt, is illustrated in Figure 2. Advisory information is presented in both the aircraft's flight data block as well as in a separate advisory window that allows controllers to accept or reject EDA recommendations. In the example shown in Figure 2, the descent-speed advisory was generated to absorb 2 minutes of required delay, evident in the ETA and STA timelines shown to the left of the display.

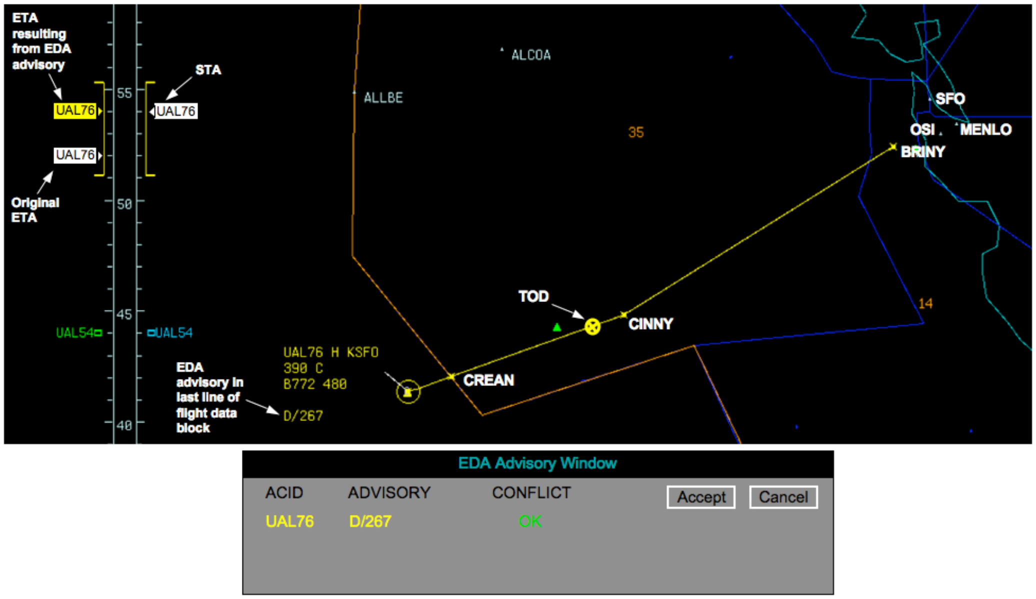

Figure 2. EDA Plan-view Graphical User Interface Showing Descent Speed Advisory

B. Profile Development and Clearance Composition

A primary challenge for OTA, as with any CDA activity, is the design of the profile constraints that define the VNAV descent profile. These constraints must be carefully chosen to allow the FMS to build a trajectory that can be flown with near-idle thrust, using only elevator inputs and normal drag-device deployment, i.e., flaps and landing gear, for vertical control. This design problem essentially reduces to one of energy management under multiple constraints. The descent trajectory must allow the potential and kinetic energy at Top-of-Descent (TOD) to be bled off at a sufficient rate to satisfy all ATC and aircraft operational constraints along the descent path, while leaving the aircraft in a suitable energy state and control configuration at the final approach fix. To further constrain the problem to allow the FMS to generate a unique idle-thrust trajectory solution for a given wind forecast, an initial descent-speed profile must be chosen. In the absence of time-based metering requirements, this speed profile is set by the FMS based on a Cost Index, computed as a ratio of time and fuel related costs for an economical idle-thrust (i.e., ECON) descent.[13] For OTA flights where meter-fix crossing times were imposed, EDA was used to override Cost Index in determining the appropriate initial descent-speed profile. This use of EDA and its subsequent impact on TOD, flight-path angle, and meter-fix arrival time is described in detail in Ref. 9.

The OTA route clearance consisted of the entire set of lateral and vertical constraints needed by the FMS for building an idle-thrust guidance trajectory. The OTA route clearance was developed iteratively, relying on extensive flight simulation with UAL and Boeing line pilots under various wind conditions and descent-speed assumptions. The primary objective was to avoid leaving the airplane low on energy relative to the VNAV path, which would trigger undesired throttle inputs from the autopilot. The second objective was to avoid leaving the airplane high on energy relative to the VNAV path, which would require speed brakes, unusual flap settings, and/or steep descent segments - all of which can increase pilot workload and passenger discomfort, while compromising desired fuel, emissions, and noise benefits.

The OTA route clearance used to support flight scenarios with EDA, including lateral waypoints assignments and associated speed/altitude crossing constraints, is shown in Figure 3. The first restriction in the descent occurred at the EDA meter fix BRINY, where the airplane was required to cross at 240 kt CAS, at an altitude of 11,000 ft. The next restriction occurred at Woodside (designated OSI), where the airplane was required to cross at 210 kt CAS, at an altitude at or above 7,000 ft. The third restriction occurred at MENLO, approximately 12 nmi from SFO, where the airplane was given no explicit speed restriction but required to cross at an altitude at or above 4,500 ft. Beyond MENLO, the remaining descent trajectory waypoints (CEPIN and AXMUL), crossing restrictions, and runway assignment (28R for these trials) were conveyed by the published approach procedure. Because the details of the approach procedure were already stored in the FMS, only the name “ILS 28R” needed to be sent via datalink along with the OTA route clearance. Upon receipt, the FMS automatically appended the detailed approach procedure to the OTA route clearance, thereby forming the basis of a complete guidance trajectory from oceanic-cruise flight to the runway.

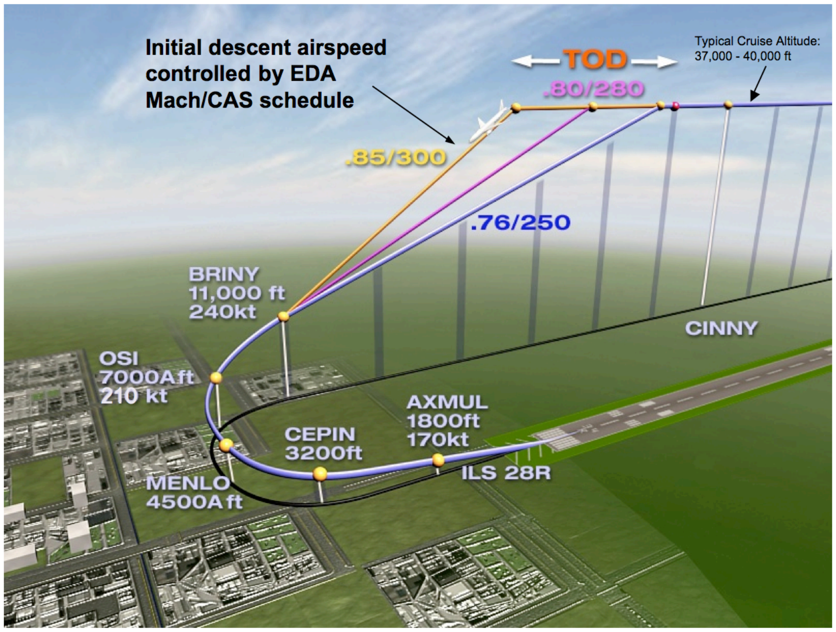

Figure 3. OTA Route Clearance Showing Profile Constraints

C. Procedures

As illustrated in Figure 4, OTA procedures were initiated in ZOA-controlled oceanic airspace, inside Oceanic Control sector 4 (OC-4). A key challenge of these flight trials was to design procedures sufficient to allow UAL76 to progress uninterrupted along the intended OTA trajectory while traversing through different regions of airspace control. Indeed, this problem represents a key obstacle to accommodating the trajectory-based operations being advocated under NextGen in today’s procedurally segmented ATC system. Today, managing flights between ATC sectors and facilities is handled through a multitude of coordination fixes, preferred routes, and airspace boundaries, together with associated speed/altitude crossing restrictions. Although these constraints help produce predictable traffic flows and divide the ATC problem into manageable, well-defined regions of control, they often occur at the expense of flight efficiency. The challenge of executing the OTA trajectory was to develop procedural techniques for allowing UAL76 to progress uninterrupted from oceanic airspace (OC-4), through domestic en route airspace (ZOA sector 35), into the Northern California TRACON (NCT) airspace, and finally to landing at SFO.

OTA procedures were initiated from the flight deck voluntarily. Flight crews were informed of the OTA trials and expected procedures through a special flight-manual bulletin distributed through UAL crew stations. Once underway, OTA procedures could be terminated at any time at the discretion of either the flight crew or ATC. The flight would then be handled by ATC using normal arrival procedures, requiring the crew to typically disengage FMS LNAV/VNAV functions. For simplicity, it was decided that once interrupted no attempt would be made to resume a CDA via OTA uplinks or procedures.

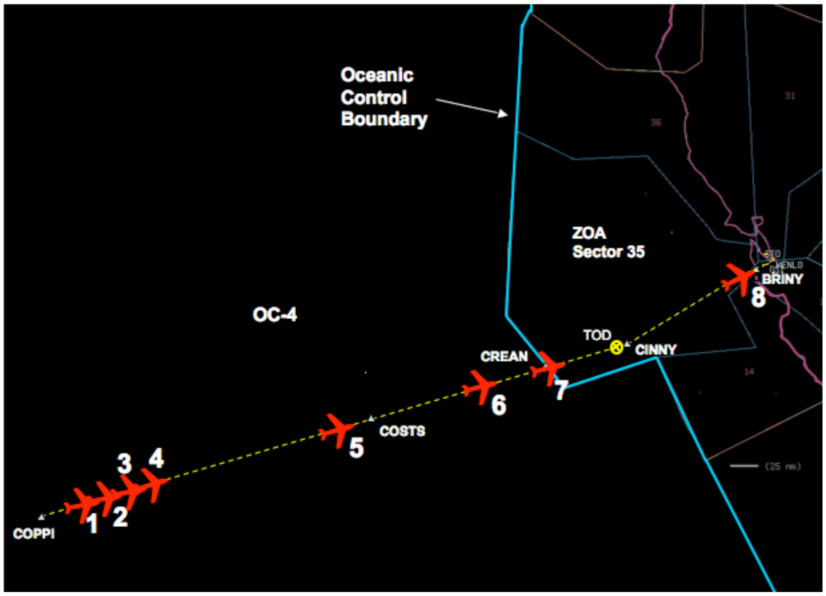

Figure 4. OTA Procedure Steps

At Step 1 in Figure 4, the flight crew requested participation around 90 min prior to their estimated landing time, approximately 700 nmi from SFO and 450 nmi prior to entering radar-controlled, domestic airspace. This request was made using a free-text CPDLC datalink message that read "Requesting OTA trials." Upon receiving the crew request, the OC-4 controller configured a new ADS-C contract within ATOP that instructed the aircraft to downlink position, weather, aircraft state, and flight-path intent data at a rate of once every 2 minutes, represented by Step 2 in Figure 4.

Following ADS configuration, the basic OTA clearance, as previously described, was up-linked to the aircraft using a FANS-loadable CPDLC message (Step 3). The message, conveyed using CPDLC Uplink Message (UM) 83, was as follows: "At COSTS cleared CREAN CINNY BRINY/N0240A110 OSI/N210A070 MENLO/A045A ILS28R." The test engineer coordinated with the ZOA-oceanic supervisor in advance to affirm the assumed landing runway. Prior to accepting the clearance, the flight crew first ensured that it could be loaded satisfactorily into the FMS to produce a continuous trajectory to the runway (Step 4). This FMS trajectory was computed in compliance with the route, speed, and altitude constraints stipulated in the OTA clearance. Upon crew acceptance of the OTA clearance, a "Wilco" message was downlinked to ATOP.

Step 5 involved the uplink of wind and temperature data to the flight deck for inclusion in FMS trajectory calculations. These atmospheric data were derived from the same RUC 2-hour forecast model used for ground-based EDA trajectory calculations. In addition to the forecast surface temperature at SFO, the data consisted of wind speed/direction at five points along the OTA trajectory corresponding to 1) cruise altitude at the waypoint CINNY, 2) cruise altitude at TOD, 3) 18,000 ft along the descent path, 4) 10,000 ft along the descent path, and 5) threshold crossing at SFO. These data were uplinked to the flight deck to provide FMS trajectory computations with the same atmospheric data available to EDA. This was essential to make valid trajectory-prediction comparisons between ground-based and airborne automation in post-trial analysis.

Step 6 involved the uplink of the EDA descent-speed advisory, intended to control arrival time at the waypoint BRINY. The advisory was obtained from a prototype EDA tool running on a laptop computer in the ZOA-Oceanic control room. Upon extracting the advisory, the test engineer relayed it to the oceanic sector controller managing UAL76. The controller then used ATOP to relay the instruction to the aircraft in a datalink message consisting of current Mach number and the advised descent CAS. Upon receipt by the flight deck, the descent speed instructions were manually entered into the FMS VNAV descent page, which resulted in a recalculation of the FMS TOD and trajectory needed to target the BRINY constraint. Once reaching TOD, the FMS commanded the airplane to initiate the descent at an airspeed equal to the current Mach number until capturing the EDA-advised descent CAS.

After leaving oceanic airspace and entering the radar-controlled domestic airspace of ZOA Sector 35, UAL76 received a voice-based pilot-discretionary descent clearance to 8,000 ft (Step 7); all clearances were given by voice from this point forward since CPDLC services are not currently available in U.S-domestic airspace. Although coordination had already taken place between ZOA and NCT at the supervisory level, the ZOA Sector 35 controller then notified the downstream, receiving controller at NCT that UAL76 was "on the OTA." Assuming allowable traffic conditions, the NCT controller, upon accepting the hand-off from Sector 35, cleared UAL76 to 4,000 ft and issued the appropriate approach clearance and runway assignment (Step 8). In general, procedures were designed so that all voice-issued altitude clearances stayed ahead of the altitude restrictions contained in the OTA route clearance being executed through the FMS to avoid interrupting the CDA.

It is recognized that these above procedures, although acceptable for limited flight trials, are likely to be too cumbersome for routine trajectory-based operations. The piecemeal voice instructions needed to communicate standard altitude clearances and crossing restrictions, while staying ahead of the OTA route clearance in the FMS, could likely be replaced by a simple "descend via the OTA" instruction. Although requiring formal changes to ATC and airline Standard Operating Procedures, this approach is currently being pursued in support of near-term deployable Tailored Arrivals.[§]

D. Test Conditions and Data Collection

The OTA trials were conducted in two phases: Phase 1 (August 17 - September 6, 2006), and Phase 2 (December 13 - January 9, 2007). The divided schedule was due to the airline's choice of when to assign a FANS-equipped B777 to UAL76. Two distinct operational test conditions were employed in the OTA trials, referred to as OTA1 and OTA2. For OTA1 flights, the initial descent speed was not stipulated by ATC; instead the pilot was free to execute a pilot-discretionary descent using the FMS ECON speed profile computed using Cost Index. OTA1 flights were conducted to help identify near-term procedural requirements and assess the immediate benefit of conducting Tailored Arrivals under light traffic conditions where EDA automation is not required. OTA2 flights, which included the BRINY metering constraint for interoperability with EDA, were used to support trajectory-prediction comparisons between air and ground automation and congested-airspace benefit assessments.

A summary of the OTA flights is shown in Table 1 with a breakdown of successfully completed events. There were a total of 40 flight opportunities over both phases of the trials. Of these, pilots volunteered to participate on 35 occasions. RUC-based winds were successfully up-linked on 27 of these occasions. The total number of uninterrupted CDA operations to the TRACON boundary and runway (independent of successful wind uplinks) were 27 and 20, respectively, of which approximately 80% were OTA2 flights. Upon taking only those flights that had successful wind uplinks together with uninterrupted CDA to the TRACON boundary, and further eliminating those with any unexplained route deviations and/or pilot-reported anomalies, 11 flights remained for the detailed trajectory analysis described later in this paper.

The quantitative data collected during the OTA trials are shown in Table 2. In addition to quantitative measurements, qualitative data were collected to refine human procedures and identify real-world issues associated with OTA deployment and execution. Qualitative observations were also used to support post-flight quantitative analysis by helping to explain any anomalies associated with a particular flight. Qualitative data gathered by project engineers included air-traffic control facility observations, jump-seat observations on the flight deck, and crew interviews at the gate upon arrival at SFO.

[§] Although not addressed here, numerous air/ground procedural findings were made pertaining to the near-term implementation of Tailored Arrivals in the current FAA system under accommodating traffic conditions.

Table 1. Flight Summary

	Total OTA Opportunities		Pilot Concurrence		Successful OTA Route Clearance Uplink		Successful Wind Uplink		Successful CDA to TRACON Boundary		Successful CDA to Landing	
Phase 1 Aug. - Sept., 2006	21		17		17		13		13		8	
	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2
	7	14	3	14	3	14	3	10	3	10	0	8
Phase 2 Dec. - Jan., 2007	19		18		18		14		14		12	
	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2
	6	13	6	12	6	12	4	10	3	11	2	10
Totals	**40**		**35**		**35**		**27**		**27**		**20**	
	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2	OTA1	OTA2
	13	**27**	**9**	**26**	**9**	**26**	**7**	**20**	**6**	**21**	**2**	**18**

Table 2. Data Collected

Source	Data Elements	Frequency
ADS-C	• Position (lat/long, altitude, time stamp) • Weather (wind speed, wind direction, temperature) • Earth reference (ground speed, vertical rate, track angle) • Air reference (Mach, heading) • Projected intent (ETAs at all downstream waypoints)	1/ (2 min)
ZOA Radar Track and Flight Plan	• Position • Ground speed, track heading	1/(12 sec)
NCT Radar Track and Flight Plan	• Position • Ground speed, track heading	1/ (5 sec)
EDA Outputs	• Trajectory predictions • Speed advisories	Event based
NOAA RUC-2	• Wind speed/direction • Temperature, Pressure	1/hr
Microphone	• Surface noise level	10 Hz
Aircraft Uplinks	• OTA route clearance • EDA speed advisory	Event based

III. Results

A. Fuel and Emissions Benefits

1. Baseline Conditions

To baseline fuel and emissions benefits, three flight-path scenarios were constructed to represent today's arrival operations under light, medium and heavy traffic conditions. These baseline scenarios were derived from observing B777 arrival traffic into SFO off CEP routes, captured for various times of day from flight-track data gathered one week prior and during the OTA Phase 1 and Phase 2 trials. The lateral and vertical track data from which baseline scenarios were derived is shown in Figure 5, color-coded for various times of day. Particularly for morning and evening flights when traffic is heaviest, these data show the inefficient lateral vectoring and altitude level-off maneuvers resulting from air-traffic control actions taken to manage separation and throughput constraints. The three baseline trajectories derived from these data are shown in Figure 6. The light-congestion baseline was used to estimate near-term benefits, since it represents traffic conditions for which OTA procedures could be invoked today without the deployment of EDA. To claim OTA benefits in comparison to the medium and heavy traffic baseline scenarios, it is assumed that EDA is available as a supporting groundside tool for developing conflict-free metering solutions.

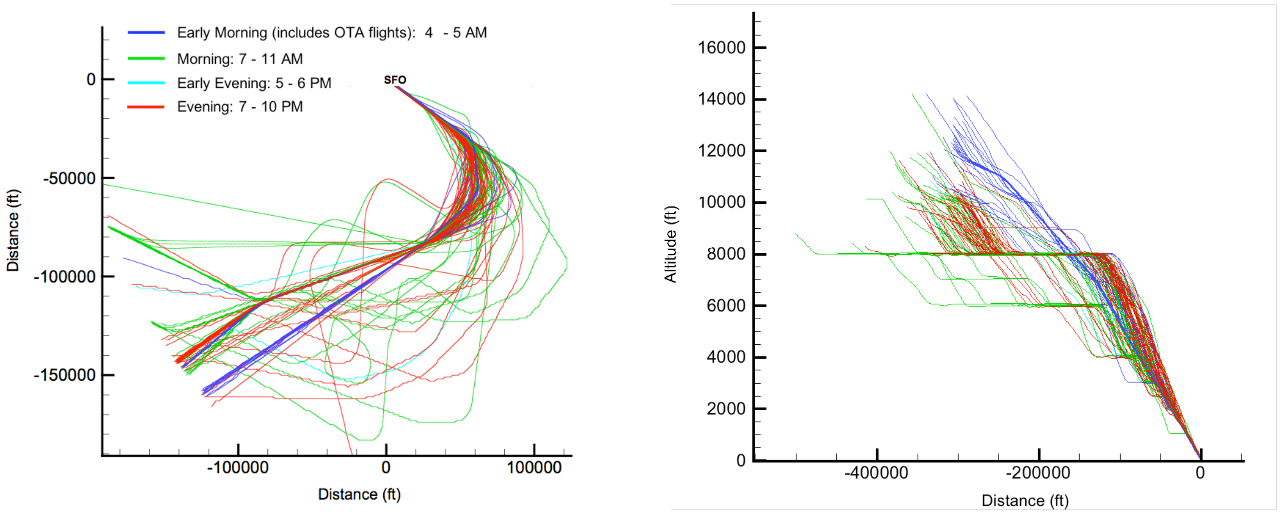

a) Lateral Tracks

b) Vertical Tracks

Figure 5. Current B777 Operations into SFO Arriving off CEP routes

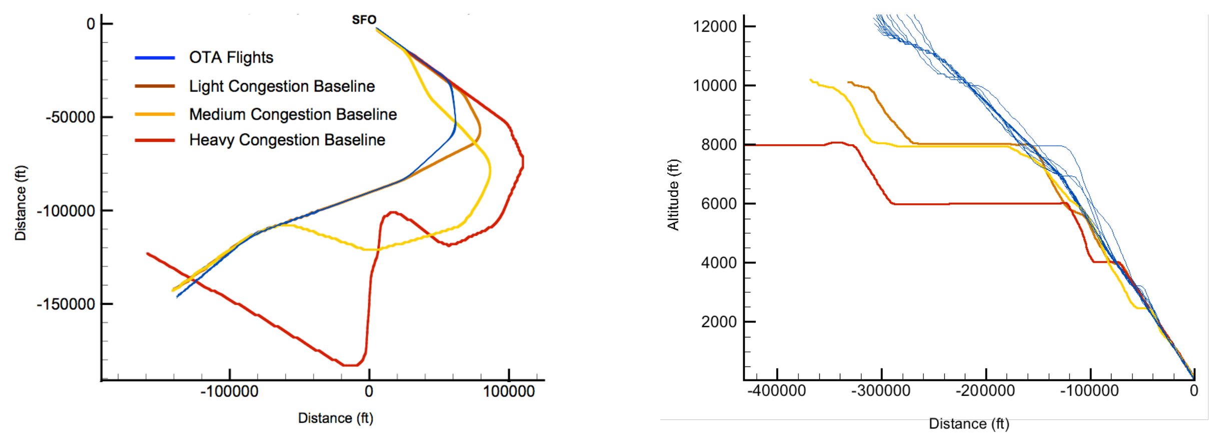

a) Lateral Baselines

a) Vertical Baselines

Figure 6. OTA Baseline Scenarios Derived from Current-Day Observations

2. *Fuel Benefits*

Due to the unavailability of direct fuel-burn measurements from the aircraft, OTA fuel benefits were estimated post flight using Boeing's proprietary BCOP/INFLT database and performance analysis software. Fuel burn associated with each of the three baseline scenarios was estimated in BCOP/INFLT for a range of initial descent CAS values: 260, 280, 300 and 320 kt. Comparative fuel burn results were computed for modeled OTA1 and OTA2 trajectory profiles using the same range of initial descent CAS values used for the baseline scenarios. For the OTA1 profile, the descent CAS was assumed constant up to the transition to the built-in FMS speed restriction of 240 kt, occurring wherever the aircraft crossed through 10,000 ft. For OTA2 flights, the descent CAS was held constant prior to transitioning to the BRINY crossing restriction of 11,000 ft and 240 kt. For each baseline and OTA scenario, Table 3 shows the distance flown and estimated fuel burn between CREAN (located on the Oceanic Control Boundary approximately 240 nmi from landing) and SFO for a B777-200. Over all baseline and OTA scenarios, the effect of initial descent CAS on total fuel burn was small, amounting to less than a 25 lb fuel difference between the optimal descent CAS (280 kt) and the least efficient descent CAS (260 kt). Therefore, to simplify the presentation of results, Table 3 shows the fuel burn averaged across all descent-CAS variations.

Under light-congestion traffic conditions, results indicate that OTA1 can provide average fuel savings of 242 lbs per flight over current-day baseline operations for B777 flights arriving SFO along CEP routes. These savings drop slightly to 227 lbs for OTA2 as a result of the additional profile constraint introduced by the meter-fix crossing restriction at BRINY. Potential fuel savings under light-congestion traffic conditions, however, are well represented by the OTA1 results, since EDA would not be required to enable CDA operations. Furthermore, since these savings are not dependent on EDA, they could be realized in the near-term using current oceanic automation systems together with trajectory-based procedures.

For medium and heavy traffic congestion scenarios, average estimated OTA fuel savings increase to 358 lbs and 3,219 lbs per flight, respectively. The dramatic increase in estimated fuel savings for the heavy traffic comparison was due to the inefficiencies inherent in the baseline scenario for achieving flow management and separation assurance under congested conditions. As seen in Figure 6, these inefficiencies included an extra 30 nmi path-stretch segment performed in level flight at low altitude (6,000 ft). To accurately estimate OTA fuel savings under heavy traffic conditions, the effect of upstream metering actions needed to approximately match the arrival times of baseline flights into SFO was considered. For these scenarios, delay absorption was required to prevent OTA flights from arriving too early, thereby violating capacity constraints. Although, ideally, some delay could be absorbed on the ground by carefully planning departure times in anticipation of CDA operations in the presence of other traffic, it is more realistic to assume that delay must be absorbed in flight, either as a result of additional EDA advisories in arrival airspace and/or regional flow-management directives applied further upstream. To compute the fuel savings associated with heavy-congestion operations, an approximately 25% fuel-burn penalty was calculated for OTA flights to account for upstream metering actions. It was assumed that this penalty was taken upstream of CREAN in cruise flight in the form of a path-stretching maneuver designed to match baseline arrival times. This penalty accounts for fuel savings in the heavy congestion scenario of Table 3 being less than the difference in fuel burned from CREAN to SFO between baseline and OTA scenarios. No such penalty was needed for OTA flights under light and medium traffic congestion, since required delay absorption could be accomplished with the descent-CAS variation already taken into account.

Table 3. Fuel Burn: OTA vs. Baseline

		Distance flown from CREAN to SFO (nmi)	Fuel burned from CREAN to SFO (lbs)	Δ Fuel: Baseline - OTA1 (lbs)	Δ Fuel: Baseline - OTA2 (lbs)
OTA	OTA1	233	7,485		
	OTA2	233	7,500		
Baseline	Light	239	7,727	**242**	**227**
	Med	244	7,858	**373**	**358**
	Heavy	273	11,680	**3,224**	**3,219**

3. Emissions Benefits

Estimates of per-flight OTA emissions reduction in comparison with baseline conditions are shown in Table 4. These results were calculated along with the fuel burn using BCOP/INFLT between CREAN and SFO. Table 4 shows the effect of OTA on the four compounds of primary concern to the environment. These are 1) carbon dioxide (CO_2) - a greenhouse gas produced as a normal product of organic-fuel combustion; 2) carbon monoxide - a poisonous gas resulting from incomplete combustion; 3) nitrogen compounds - primarily nitric oxide (NO) and nitrogen dioxide (NO_2) resulting from high temperature combustion and commonly associated with ozone and smog; and 4) all species of hydrocarbons (C_xH_x) - volatile organic compounds resulting from unburned or partially burned fuel passing through the engine.[14] As shown in Table 4, results suggest that an idle-thrust, OTA1 descent to the runway can reduce CO_2 emissions by 761 lbs per flight in comparison with B777 operations on similar routes conducted during light traffic congestion. In comparison with medium and heavy traffic-congestion baselines, OTA2 has the potential to reduce CO_2 emissions by as much as 1,128 lbs and 10,137 lbs per flight, respectively. These large greenhouse gas reductions reflect the approximate 3:1 ratio between fuel burned and CO_2 emitted, resulting from the basic combustion chemistry of jet fuel.[14]

Table 4. Emissions: OTA vs. Baseline

		Emissions from CREAN to SFO					
		CO (lbs)	NO_x (lbs)	C_xH_x (lbs)	CO_2 (lbs)	ΔCO_2: Baseline - OTA1 (lbs)	ΔCO_2: Baseline - OTA2 (lbs)
OTA	OTA1	17.3	133.9	0.77	23,572		
	OTA2	17.5	134.0	0.77	23,618		
Baseline	Light	19.1	135.6	0.82	24,333	**761**	**715**
	Med	19.3	138.1	0.84	24,746	**1,174**	**1,128**
	Heavy	19.7	238.6	0.95	36,782	**10,152**	**10,137**

B. Trajectory Performance

Data showing the accuracy of FMS arrival-time predictions to the meter fix (BRINY) in comparison with the actual BRINY crossing time, as a function of time-to-go to BRINY, are shown in Figure 7. These data show the expected, continual improvement of airborne predictions as the airplane progresses through oceanic and domestic airspace. Of note are the distinct improvements in prediction accuracy for several flights that can be observed to correspond to the wind and descent-speed uplink events. The latter, of course, was expected whenever the airplane's original VNAV descent speed, based on Cost Index, differed from the EDA descent speed clearance that was uplinked and executed.

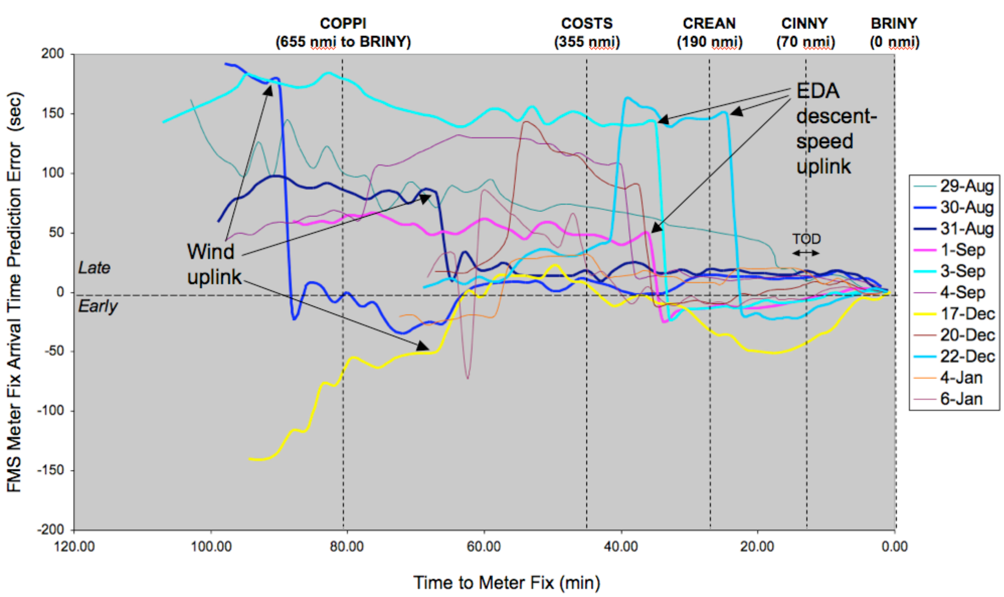

Figure 7. Progression of FMS Arrival-Time Predictions to Meter Fix

EDA and FMS arrival-time estimates to the meter fix, resulting from trajectory predictions generated 200 nmi/ 25 min upstream at CREAN, are illustrated by the histogram in Figure 8. These results pertain to a single prediction at CREAN, i.e., without updates as the flight progresses. These data show mean arrival-time prediction accuracies of 2 seconds early and 3 seconds late for FMS and EDA, respectively, with a similar dispersion about the mean ($\sigma \approx$ 20 sec) for both. These data suggest that airborne and ground-based automation can predict arrival times over a 25 min horizon with similar accuracy and precision, assuming shared wind and descent-speed-intent information.

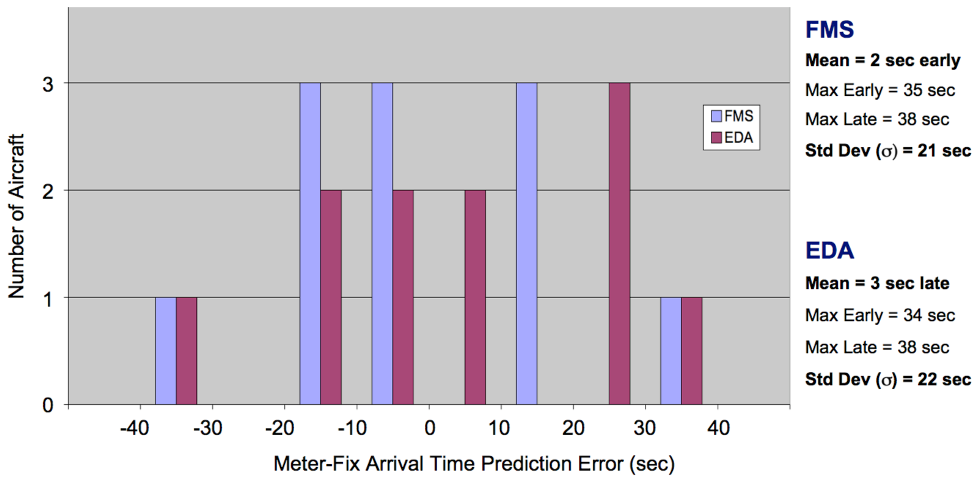

Figure 8. Histogram of EDA and FMS Arrival-Time Prediction Error to Meter Fix (200 nmi/ 25 min horizon)

Because EDA conflict avoidance functions require accuracy at each point along the trajectory prediction, not just at the meter fix itself, the entire EDA prediction was compared to surveillance truth data. The overall accuracy of EDA trajectory predictions for several look-ahead times is shown in Table 6. These results capture the error in altitude, along-track position, and cross-track position along the entire trajectory prediction to the meter fix for time horizons of 23 min, 20 min, and 17 min. These look-ahead times were chosen so that predictions for all flights could be initiated between the Oceanic Control Boundary and TOD to allow use of ARTCC radar surveillance as truth data, captured every 12 seconds. Inherent latencies associated with the radar data (a constant bias for each flight ranging from 6 to 12 seconds) were identified and removed in post processing by calibrating them with the time-stamped airplane position reports received via ADS-C. Figures 10 and 11 show the altitude and along-track error for all flights as a function of time for a 23-minute prediction time horizon. A similar comparison of FMS trajectory prediction error as a function of time was not possible, since only the FMS time estimates to downstream waypoints were available, not the full trajectory predictions upon which those estimates were based.

Results in Table 6 show that EDA's mean altitude prediction error, ranging between -500 and -700 ft, does not vary substantially with time horizon. This is because most altitude error occurs in the descent phase of flight, which is fully contained within each prediction horizon and influenced primarily by speed intent and modeled aircraft performance characteristics rather than initial conditions in cruise. The negative mean altitude error is due to EDA predicting the airplane to reach the meter-fix crossing altitude earlier than what actually occurred due to EDA's modeling of the deceleration to the meter-fix crossing speed (240 kt) using a level flight segment. Although this assumption works well for the current-day TMA operations that EDA was originally designed to support, it is a poor model for CDA operations specifically designed to avoid level flight segments. The result is a ground-based prediction to the meter fix that lags the actual flight operation. This results in an earlier TOD estimate (evident by the initial spike in altitude error in Figure 10) and lower overall altitude prediction in comparison to truth. The same phenomenon results in the negative mean along-track prediction errors (ranging between -0.6 and -1.3 nmi) seen in Table 6. Unlike altitude errors, along-track errors averaged over the entire trajectory prediction grow significantly as the time horizon increases from 17 to 23 minutes. This is due to groundspeed prediction error – resulting primarily from remaining wind uncertainty – occurring over both cruise and descent.

This EDA-prediction analysis includes the effect of sharing wind data and descent-speed intent with the airplane in actual operations. To see the importance of shared descent-speed intent, in particular, on air/ground predictions, a simple study was carried out looking at EDA trajectories under descent-speed-intent assumptions ranging from 250 kt to 320 kt CAS. Over a 25-minute prediction horizon, the altitude and along-track differences at any given time can be as large as 6,500 ft and 20 nmi as a result of descent-speed-intent uncertainty.

Table 6. EDA Trajectory Prediction Accuracy for 23 min, 20 min, and 17 min Time Horizons

	23 min prediction				20 min prediction				17 min prediction			
	Mean	Std Dev	Min	Max	Mean	Std Dev	Min	Max	Mean	Std Dev	Min	Max
Altitude Error (ft)	-502	714	-3301	809	-664	696	-3290	564	-693	670	-3473	479
Along-Track Error (nmi)	-1.33	1.45	-4.80	2.30	-0.73	0.99	-2.82	1.98	-0.61	0.78	-2.44	1.85
Cross-Track Error (nmi)	-0.12	0.23	-1.19	0.37	-0.11	0.21	-1.09	0.15	0.09	0.21	-1.11	0.24

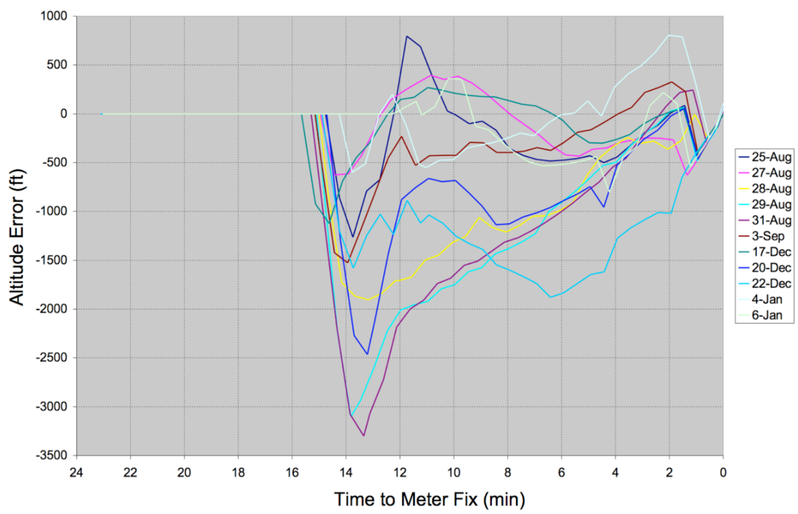

Figure 10. EDA Altitude Prediction Error (23-min time horizon)

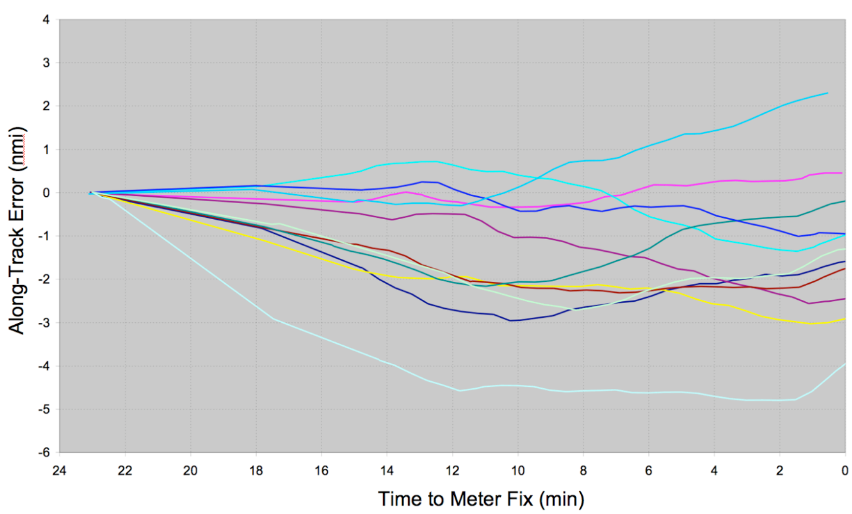

Figure 11. EDA Along-Track Prediction Error (23-min time horizon)

IV. Conclusions

The San Francisco Oceanic Tailored Arrivals field trials demonstrated the ability to conduct highly efficient Continuous Descent Approach operations under real-world conditions with commercial airplanes. By integrating advanced air and ground automation over datalink in oceanic airspace, these trials provide a step towards understanding the feasibility and benefits of conducting trajectory-based arrival operations under NextGen. To progress towards enabling Tailored Arrivals under congested traffic conditions where benefits are greatest, NASA's ground-based EDA automation was used to tailor trajectory solutions to accommodate time-based metering constraints. Data gathered during these trials was used to perform an assessment of EDA trajectory prediction accuracy and precision in comparison with FMS predictions and surveillance truth data. Results show that trajectory prediction errors can be greatly reduced through the uplink of wind and descent-speed-intent data, and that similar arrival-time prediction performance can be achieved between air and ground automation. These results are important because accurate and compatible prediction performance between air and ground automation is fundamental to both the planning and execution of trajectory-based arrival operations.

Finally, an initial benefits assessment based on the real-world data gathered during these trials shows substantial per-flight reductions in fuel burn and environmental emissions afforded by Tailored Arrivals, especially in comparison with baseline arrival operations conducted under heavy traffic conditions. These results are particularly compelling in the presence of today's record high fuel costs and increased environmental awareness.

Acknowledgments

The authors would like to acknowledge the substantial contributions of FAA personnel from Oakland Center and Northern California TRACON in planning and executing these operational field trials. In addition, the support received from FAA headquarters (ATO-E) was instrumental in approving and coordinating this activity. Critical systems engineering, data analysis, and human-factors contributions were made by project personnel from Boeing, Sensis Corporation, the San Francisco Noise Abatement Office, Lockheed Martin, NASA Code TH, QSS Group Inc., University of California Santa Cruz, and San Jose State University. Last but not least, the authors would like to thank to the subject-matter experts and pilot participants of United Airlines that enabled these trials to take place.

References

[1]Gouldner, D., Boerner, S., "Evaluating NGATS Research Priorities at JPDO," *Proceedings of the 6th AIAA ATIO Conference,* Wichita, KS, Sept. 2006

[2]Prevot, T., Callentine, T., Lee, P., et al., "Co-Operative Air Traffic Management: Concept and Transition," *Proceedings of the AIAA Guidance, Navigation and Control Conference,* San Francisco, CA, Aug. 2005.

[3]Anderson, L.R. & Warren, A.W., "Development of an Advanced Continuous Descent Concept Based on a 737 Simulator. *Proceedings of the 21st Digital Avionics System Conference, AIAA*, Irvine CA. Oct., 2002

[4]Watt, J., Follet, J., Mead, J., et al., "In Service Demonstration of Advanced Arrival Techniques at Schiphol Airport," *Proceedings of the 6th AIAA ATIO Conference*, Wichita, KS, Sept. 2006

[5]Elmer, K., Wat, J., Warren, A., Tong, K., et al., "Community Noise Reduction Using Continuous Descent Approach: A Demonstration Flight Test at Louisville," *Proceedings of the 9th AIAA Aeroacoustics Conference*, Hiltin Head, SC, May 2003

[6]Clarke, J., Ho, N., Ren, L., Brown J., et al., "Continuous Descent Approach: Design and Flight Test for Louisville International Airport," *Journal of Aircraft*, Vol. 41, No. 5, 2004, pp. 1054-1066

[7]Reynolds, T., Ren, L., Clarke J., et al., "History, Development and Analysis of Noise Abatement Arrival Procedures for UK Airports," *Proceedings of the 5th AIAA ATIO Conference*, Arlington, VA, Sept. 2005

[8]Roberts, C., Cornell, B., Mead, R., "Tailored Arrival Joint Project Phase One, Final Report," Dec. 2004

[9]Coppenbarger. R., Lanier, R., Sweet, D., et al., "Design and Development of the En-Route Descent Advisor for Conflict Free Arrival Metering," *Proceedings of the AIAA Guidance, Navigation and Control Conference,* Providence, RI, Aug. 2004.

[10]Green, S. "Descent Advisor Preliminary Field Test," *Proceedings of the AIAA Guidance, Navigation and Control Conference*, Baltimore, MD, Aug. 1995

[11]Farley, T., Foster, J., Hoang, T., and Lee, K., "A Time-Based Approach to Metering Arrival Traffic to Philadelphia," *Proceedings of the 1st AIAA ATIO Conference*, Los Angeles, CA, Sept. 2001

[12]Grabbe, S., Sridhar, B., "Central East Pacific Flight Routing," *Proceedings of the AIAA Guidance, Navigation and Control Conference*, Keystone, CO, Aug. 2006

[13]Chakravarty, A., "Selection of an Optimal Cost Index for Airline Hub Operation," *Journal of Guidance, Control, and Dynamics,* Vol. 8, No. 6,1985, pp. 777-781

[14]Aviation & Emissions - A Primer, FAA Office of Environment and Energy, Jan. 2005