Eleventh USA/Europe Air Traffic Management Research and Development Seminar (ATM2015)

Performance Improvements Through Trajectory Feedback in the Future Collaborative Flight Planning Environment

Stéphane Mondoloni., Ph.D, Sheng Liu, Daniel Kirk, Ph.D.
The MITRE Corporation
Center for Advanced Aviation System Development
McLean, USA
smondoloni@mitre.org, sliu@mitre.org, dkirk@mitre.org

***Abstract*— Pre-departure feedback on trajectory constraints will be enabled through future flight planning provisions under development by ICAO. Select operational benefits of such feedback were investigated. Two benefit mechanism are described and quantified through the use of operational data to drive Monte Carlo simulations. One mechanism involves the use of pre-departure feedback to enable an airspace user to select the route that is lowest cost when considering ATC constraints. This mechanism yields an average 0.2% improvement in time and fuel costs. Individual flights may gain up to 2% improvements. The benefit is highly dependent on the flight's origin and destination. A second benefit mechanism allows an airspace user to pre-emptively take ground delay to reduce the likely airborne delay on arrival. This is accomplished through feedback on the estimated delay distribution based upon shared information on all relevant flights. A simulation of arrivals into London Heathrow as an example of the benefit mechanism revealed the potential to shift an average of 2.3 minutes of airborne holding to the ground for 138 flights per day. This represent an annual gain on the order of 10 million kilograms of fuel for that one airport alone.**

Keywords-Flight planning, Trajectory, Collaboration, FIXM

I. Introduction

In 2004, the International Civil Aviation Organization (ICAO) established the Flight Plan Study Group (FPLSG) "for the purpose of developing a proposal for revision of the flight plan provisions, including the flight plan form and associated operating practices [1]" The FPLSG decided to split the task into two sets of changes: a near-term amendment to the Procedures for Air Navigation Services – ATM [2] for global implementation in November of 2012, and a more far-reaching overhaul of the flight planning process which became the Flight and Flow Information for a Collaborative Environment (FF-ICE) Concept document [3]. This far-reaching concept was fully supported at a global level by the 12th Air Navigation Conference in November of 2012 [4].

Some of the key changes described by the FF-ICE Concept include: 1) a significantly more flexible process for defining the information items that can be exchanged, 2) the use of the Flight Information Exchange Model (FIXM) as the information exchange standard, 3) a more collaborative and dynamic flight planning process including the exchange of four-dimensional trajectory (4DT) information, and 4) support for a transitional, mixed-mode environment.

Given the extent of changes envisaged by the FF-ICE Concept, a phased implementation is anticipated and described through the FICE thread of the ICAO Aviation System Block Upgrades (ASBU) [5]. Provisions applicable to the first block, ready for implementation in the 2018-2023 timeframe, are presently being developed by ICAO and referred to as FF-ICE Step 1.

Step 1 includes the pre-departure exchange and negotiation of the 4DT between the ATM Service Provider (ASP) and the Airspace User (AU). While the scope of Step 1 is described at a high-level in the ABSU description, refinement of the scope and operational scenarios describing the exchange and use of the 4DT have been developed collaboratively between the Federal Aviation Administration (FAA) and the SESAR Joint Undertaking (SJU). These have been presented at ICAO and accepted to form the basis for further development of FF-ICE Step 1 flight planning provisions.

This paper first summarizes the use of the 4DT envisaged for Step 1 and describes specific benefit mechanisms derived from feedback on the 4DT. Two areas are focused on: 1) the ability of the AU to select a more optimal plan given feedback on the constraints and anticipated ATC changes to a proposed 4DT, and 2) the ability of the AU and a downstream ASP to engage in flow management enabling the shifting of airborne delay to ground delay across many Flight Information Region (FIR) boundaries.

II. The 4DT in FF-ICE Step 1

Airspace Users complying with FF-ICE Step 1 provisions, who operate a portion of their flights within airspace controlled by a compliant ASP will be able to engage in pre-departure

Sponsored by the Federal Aviation Administration, NextGen Office (ANG)

trajectory negotiation. The process for such negotiation involves:

- The ASP makes available some known constraints within their airspace
- The AU develops a flight plan, together with a 4DT meeting their operational objectives and the above constraints
- The AU provides the flight plan and 4DT to FF-ICE Step 1 compliant ASPs along the route
- The ASPs evaluate the flight plan and 4DT and provide feedback described below
- The AU incorporates the information provided in the feedback from one or more ASPs into a revision of the flight plan and 4DT to better meet constraints and objectives
- New feedback may be provided by ASPs as circumstances change
- When multiple ASPs are involved in Step 1, each ASP provides feedback within their own airspace to be reconciled by the AU

Feedback is provided to the AU expressing any flight plan amendments and constraints the ASP expects, given what is presently known. This feedback includes the application of re-routes, and speed and altitude restrictions. These are reflected in a trajectory-route pair [6] provided back to the AU. If known, constraints such as miles-in-trail may also be provided (although their impact may not be reflected in the 4DT). Through the use of such feedback, an AU may pick a flight path resulting in a more optimal outcome.

When developing a flight plan, an AU may wish to receive feedback on multiple candidate plans to facilitate the choice of the best plan. In such a case, a collection of trajectory-route pairs may be provided if so enabled by the ASP.

Throughout the above exchange, information is provided using FIXM version 4.0. However, locally-applicable feedback may also be provided through FIXM extensions. One key aspect of FIXM is the ability of an ASP or a region to implement local extensions. These extensions contain information items that are not described within the FIXM global core (i.e., for items that have global applicability). Extensions may be used for local procedures, or for investigation of new capabilities. One example of such a new capability is the use of a probabilistic evaluation of airborne delay to allow AU-initiated ground delay in one ASP to mitigate further airborne delay or fuel cost in another.

The above benefit mechanisms are described further in subsections A and B below.

Additional benefit mechanisms have previously been reported through the use of operator-provided data to improve operations using trajectory prediction. FF-ICE Step 1 and FIXM enable the exchange of such information, such as aircraft weight and operator-provided climb and descent profiles. These improve flow management through improved demand prediction (e.g., [7-9]), improve time-based metering through improved arrival time estimation (e.g., [10-11]), and improve conflict detection and resolution through improved efficiency owing to reduced false and missed alerts (e.g., [12-15]). These benefit mechanisms are not re-evaluated herein.

A. *Use of 4DT Feedback for Route Optimization*

Under present operations, when a flight plan is filed, time and fuel computations for that filed plan are based upon a trajectory within the airspace user's flight planning system. The ATC system subsequently applies rules which may result in a modification to the routing, and additional speed or altitude constraints.

With experience, an AU may be able to anticipate the routing and constraints that the flight will face. However, these are not static and are the result of the application of elaborate rules. Through the provision of feedback on a filed 4DT, the AU can be informed of the known changes to a flight prior to committing to a flight plan. This allows the airspace user to select the flight plan resulting in an optimal solution.

We consider an airspace user with a choice of N trajectories, each with some combined fuel and time cost C.

$$C_i = c_i + \varepsilon_i + \eta_i \quad where\ i \in [1, N] \qquad (1)$$

Prior to departure, the expected total cost of a 4DT (C_i) is a random variable expressed as the sum of:

- A deterministic nominal cost (c_i) representing the filed 4DT
- The cost of perturbations to the 4DT that are known by automation (ε_i), but are presently not shared
- The cost of other perturbations to the 4DT that occur when the flight operates (η_i) such as tactical vectors

Without feedback, optimization on C is based upon the expected value of both types of perturbations.

$$C^* = \min_i \{c_i + E(\varepsilon_i + \eta_i)\} \qquad (2)$$

With feedback, perturbations known by automation (ε) are shared with the AU. The actual value can be used rather than the expected value. In addition, the expected value of unknown perturbations (η) can be conditioned on feedback, such as whether the flight is subject to miles-in-trail restriction.

$$C^* = \min_i \{c_i + \varepsilon_i + E(\eta_i | feedback)\} \qquad (3)$$

The trajectory feedback mechanism provides improved knowledge of the expected cost of a proposed flight plan. In principle, this allows a more optimal solution to be reached. This paper investigates and quantifies the benefit of feedback on the perturbations known to automation (ε) using operational data across select city pairs.

B. *Use of 4DT Feedback for Ground/Air Delay Optimization*

Where possible, it is beneficial for flights to incur delays on the ground versus when airborne. This is primarily due to lower ground vs. airborne direct operating costs. Through provision of a 4DT pre-departure, update of this 4DT en route, combined with knowledge of the uncertainty in these predictions, a distribution of demand can be obtained for any capacitated resource (e.g., an airport, airspace, oceanic tracks). This distribution can be used to determine the distribution of airborne delay expected at the capacitated resource.

With knowledge of the airborne delay distribution, an airspace user might take a ground delay to minimize the expected value of the total delay cost. Greater benefits are expected through agreements with the ASP responsible for allocating the airborne delay. For example, agreements could provide credit for ground delay taken elsewhere.

This benefit is quantified more fully below through a Monte Carlo simulation driven by statistics from operational data.

III. APPROACH

The benefit mechanisms pertaining to feedback on the 4DT were quantified using operational, modeled and simulated data. The data used and methodology is described below.

A. Approach for 4DT Feedback for Route Optimization

Select city pairs in the continental United States were evaluated across a total of 15 days in 2014. The city pairs consisted of a mix of stage lengths from 1000 to 3680 km with flows in directions both with and against forecast winds. Flights operating between the city pairs had two trajectories computed and evaluated for fuel and time: 1) a "Desired 4DT" representing the 4DT corresponding to the last filed flight plan, and 2) an "ATC Intended 4DT" corresponding to the 4DT as constrained through the application of ATC Intended routes and restrictions. These ATC Intended routes and restrictions, together with the rules for their application, are not explicitly made available to the AU today.

The process is illustrated in Figure 1. The Joint En Route Decision Support System Infrastructure (JEDI) environment (see [15-16]) converted the route in the filed flight plan, applied ATC Intended routes when required and eligible, and applied restrictions to the flights. The JEDI environment used adaptation data consistent with that applied by National Airspace System (NAS) En Route automation. While JEDI computed a 4DT, a separate application using Base of Aircraft Data (BADA) was required to obtain the fuel and time consistent with flight-specific parameters. These parameters not known from the flight plan, such as weight, were assigned to a flight consistent with the methodology used in [14].

While the 4DTs provide an estimate of the fuel and time, they do not express the relative value of the two measures. To obtain this relative value, a Cost Index (CI) [17] for the filed flight was estimated by perturbing the speed on the desired 4DT and obtaining resulting fuel and time perturbations. If the original filed plan is an optimal based on CI, the Cost Index is simply:

$$CI = -\frac{\Delta\text{Fuel}}{\Delta\text{Time}}\left(\frac{1}{100}\right) \qquad (4)$$

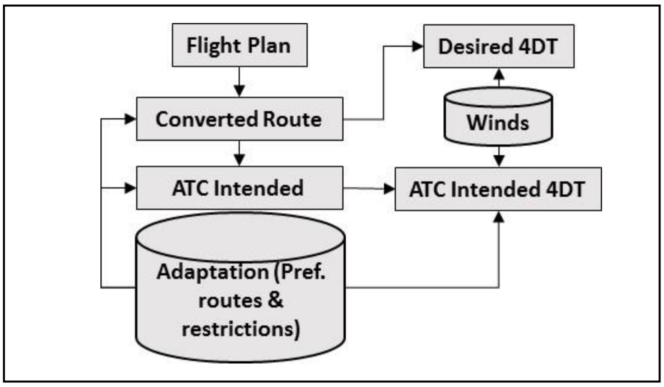

Figure 1. Process for computing Desired and ATC Intended 4DT

It is recognized that the above strictly provides an estimate for the purposes of investigating potential benefits herein. The above CI is used to combine the fuel and time costs into an estimated cost for each 4DT computed for a flight.

Each flight provides a single instance of the variables (c_i, ε_i, η_i) in (1) for a specific route choice i. As multiple flights may operate on a given filed route (i) across multiple days, the distribution of cost perturbations (ε_i, η_i) are obtained for that route. These perturbations are normalized by the base cost (c_i) to remove the effect of aircraft size and efficiency. Since perturbations typically add flight distance and additional fuel/time due to restrictions, as a first-order approximation, these are proportional to total cost.

The above is repeated across all routes ($i \in [1,N]$) observed between city pairs to gather statistics for each route choice option. These statistics are used as input to (2) and (3) to obtain optimal flight paths representing optimization without and with 4DT feedback respectively. The benefits are computed through Monte Carlo runs of the perturbations imposed on top of the nominal costs. Each instance samples (ε_i) from its distribution; however when feedback is not available, ε_i is not known when selecting the route choice. The difference between the realized costs using the optimization in (2) versus (3) provides the benefit of the 4DT feedback.

In this study, we did not investigate the conditioning of the η_i distribution based on additional feedback. Additional benefits are expected from such conditioning.

B. Approach for Ground/Air Delay Optimization

The evaluation of this benefit mechanism used operational data to obtain flight statistics which drove a multi-layered Monte Carlo simulation. As a constrained resource, we looked at flights inbound into London Heathrow (LHR) as data was available through the U.S. Traffic Flow Management System (TFMS) on both departure and arrival flights. Arrival flight data is available as a result of international data-sharing agreements. A simulation was developed to optimally trade airborne holding against ground delay using probabilistic estimates of delay. In order to develop a probabilistic model of arrival demand for that simulation, estimates of departure and en-route uncertainty were obtained from operational data.

While LHR was chosen for this simulation, it is fully recognized that airspace constraints may preclude the extraction of this benefit mechanism at this location; however, the analysis provides an indication of the benefit magnitude when such an approach is applied to other constrained resources.

Flight planning data was obtained for flights departing the US towards LHR. A total of 15 winter and 15 summer days was investigated. An estimate of the accuracy of the pre-departure Estimated Elapsed Time (EET) was obtained by comparing the observed flight time to the EET provided in the flight plan. An initial comparison indicated an average error (bias) of 10.8 minutes. This bias was due to in part to holding, vectors on arrival and effect of arrival airport configuration. When those effects are removed on a flight-specific basis using track data, the EET error bias is -1.7 minutes. The distributions of uncertainty for both cases are shown in Figure 2. Corrected data has a standard deviation of 10.3 minutes.

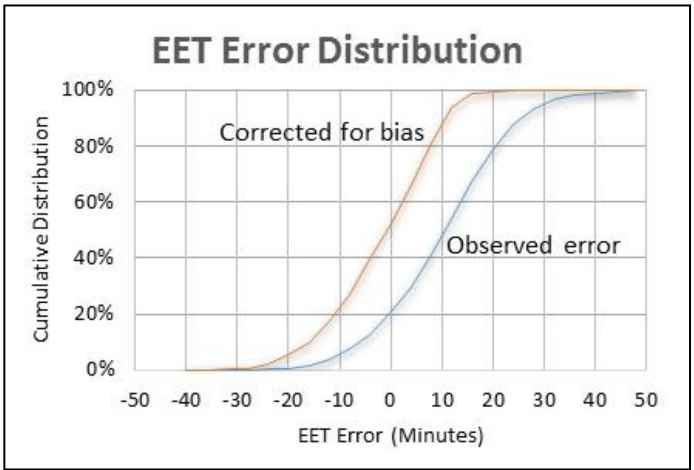

Figure 2. Distribution of EET error for U.S. flights bound to LHR

It is expected that any predicted 4DT would have the same or better accuracy on time than the EET in the present flight plan. One item under consideration for FIXM 4.0 is the inclusion of trajectory-specific temporal uncertainty as part of the 4DT. This investigation represents an estimate using uncertainty not conditioned on flight-specific data quality.

Departure delays incorporate the sum of pushback delays, taxi times and wait times at the queue. These delays can vary significantly with departure airport. Data into LHR was obtained for 6 months by comparing actual recorded departure times with planned departure times. These data were used to obtain distributions of departure delays for London-bound flights from other sources. A Gamma distribution was used to fit the departure delay distribution by airport with a shifted origin to capture instances of early departures (e.g., see Figure 3).

Figure 3 illustrates that departure delay (d), normalized through the mean delay and standard deviation (σ) can be estimated through a Gamma distribution across a selection of airports in North America, Asia and Europe. A few airports with tightly controlled departure times exhibited very narrow distributions. For other airports, data quality was an issue leading to unreasonable distributions. For the purposes of simulation, departure delay was not conditioned on departure airport. In this case, a mean of 27 minutes and standard deviation of 6 minutes was used as a representative case using the average mean and standard deviation. These were obtained through a six-month average across all flights into LHR, some outliers were discarded.

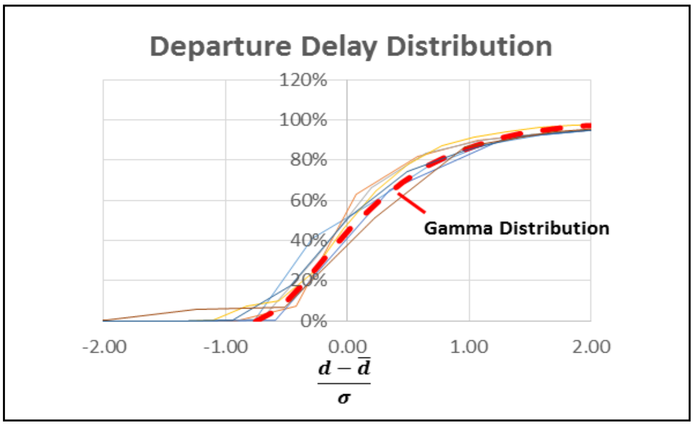

Figure 3. Distribution of actual versus fit departure delays (multiple airports)

1) Simulation Approach

A Monte-Carlo simulation is developed in which two methods are applied for dealing with arrival capacity constraints: 1) constraints are dealt with tactically through the use of airborne delays on arrival, or 2) each flight is provided with a distribution of expected arrival delays and provided with an opportunity to mitigate these through voluntary departure delays.

In the first case, each flight has an initial arrival time estimated by adding the departure delay, nominal planned flight time, and EET error to a scheduled departure time.

$$\hat{t}_{arr} = t_{sched} + \varepsilon_{dep} + t_{flight} + \varepsilon_{EET} \quad (5)$$

The departure delay and EET errors are sampled from distributions previously determined. The scheduled time and nominal planned flight time is obtained from operational data. The initial arrival time of all flights across a single day is used to allocate airborne arrival delay on a first-come first-served (FCFS) basis. The process is repeated 100 times to capture statistics on airborne delay. This case provides an estimate of the airborne delay when all delay is taken tactically.

For the second case, two nested Monte Carlo processes are run. The outer process represents the events of a single simulated day. The inner process is used to obtain a distribution of expected arrival delay, given information known at departure. The approach is best explained through an example.

1. An estimate of all flights' departure time is first obtained by adding to the scheduled time a sampled departure delay. This provides an estimate of all

flights' departure time on one simulated day. Each flight is then processed sequentially by this estimated departure time.

2. When a flight is ready for departure, an inner Monte Carlo simulation is executed for that flight to obtain a distribution of expected airborne delays, given the estimated departure time. This distribution is obtained as follows:

    a. Flights which have not yet departed have their arrival time estimated in accordance with (5). Since the departure time is not known for flights not yet airborne, this error is independent of the actual departure time.

    b. All flights have an EET error assigned to them and sampled for the inner process. The EET variance was assumed to be proportional to the EET. Flights which are airborne are assumed to have flown some distance with uncertainty. The variance in the remaining flight time EET error is reduced in proportion to time-to-go.

    c. Using the initial arrival times from (a) and (b) above, airborne flights which have taken a ground delay are provided credit for their ground delay. This re-orders the flights. The arrival time allocation process is applied to the re-ordered flights to obtain the estimated airborne delay. Regardless of the credit provided, flights cannot be assigned an arrival time prior to the expected arrival.

    d. The estimated airborne delay is approximated through multiple iterations of the inner process. This provides a distribution of expected airborne delays for the single flight being evaluated.

    e. The single flight may take a ground delay to mitigate a probability of receiving an airborne delay and then departs and is incorporated into the next flight's demand. The ground delay is selected through a threshold on the probability of airborne delay.

3. The above process repeats for all departing flights.

4. An initial arrival time is then computed for all flights, considering the EET uncertainty and consistent with the outer process times en route. In the same manner as (2c) above, the arrival airborne delay is estimate taking into account a credit for departure ground delay (see Figure 4).

Figure 4 illustrates a postulated credit assignment process. Flights 123, 334, etc. through 181 reach terminal airspace ready for arrival at an estimated time. Flights are slotted for landing based upon available arrival capacity and are assigned an arrival time based upon their arrival order. The difference between the time the flight is ready for arrival and the assigned arrival time is the airborne delay that must be taken by that flight (illustrated for flight 237). Flights that took a ground delay are provided a credit which shifts the ranking used to compute the slot allocation (e.g., flight 157).

The process for assigning credit for ground delay meets several important criteria. First, when no ground delay is taken by any flight, the process is identical to the scenario without ground delay. Second, in a fully deterministic scenario in which all flights take delay on the ground, the order and total delay value assigned to each flight is preserved.

There are two possible approaches to computing the expected delay distribution: 1) the AU may use real-time flight data publication to compute this delay, or 2) the ASP may provide this distribution as feedback. In the second case, new items would have to be incorporated in FIXM.

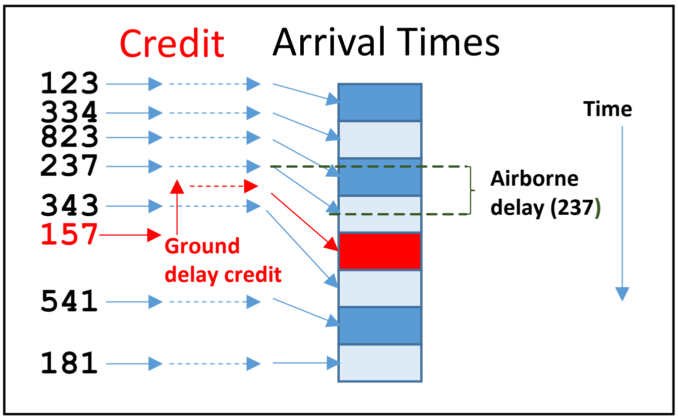

Figure 4. Example providing credit for ground delay prior to arrival delay assignment. Flights cannot be scheduled to arrive earlier than projected (not shown). Earlier flights are at the top.

IV. RESULTS

Simulations as described previously were run for both key benefit mechanisms with results described below.

A. *Results for 4DT Feedback for Route Optimization*

Fifteen days were evaluated in 2014 for flights operating between 5 city-pairs. A total of 1010 flights were evaluated across 9834 routing options to estimate the effect of feedback on the quality of the optimal solution.

For each flight evaluated, a Cost Index was inferred based upon the computed trajectory. The distribution of Cost indices is shown in Figure 5 and falls within range of values typically used by transport aircraft. These flight-specific cost indices were used to compute the flight-specific costs compared below.

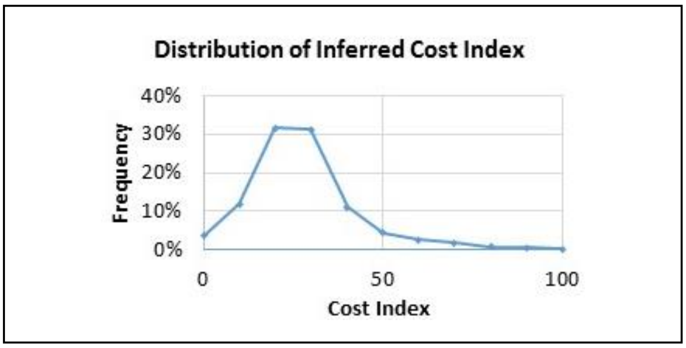

Figure 5. Distribution of inferred Cost Index

A comparison of the cost of the trajectory corresponding to the filed route without constraints to the trajectory based on the ATC Intended route with constraints is shown in Figure 6. This figure shows the average percentage cost increase together with one standard deviation and the maximum difference observed. Across all flights, an average of 0.6% and standard deviation of 0.75% was observed. This value is significantly dependent on city pair as illustrated in Figure 6. When comparing only the fuel component of total cost, an average increase in 0.7% and standard deviation of 1.1% was observed.

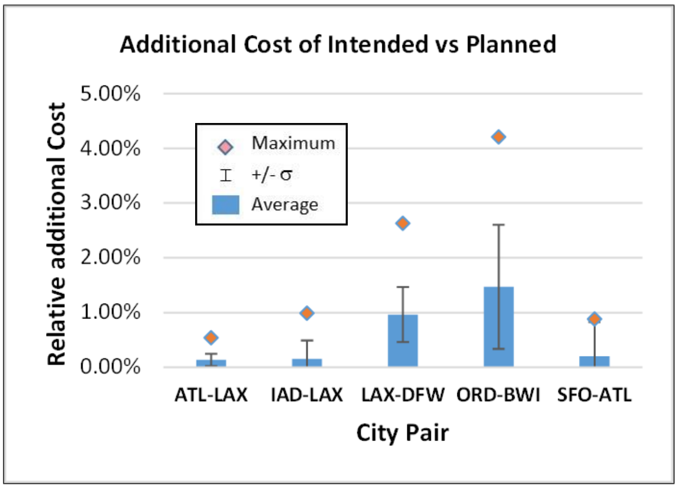

Figure 6. Effect of ATC Intended routing and constraints on cost by city-pair

With FF-ICE Step 1, feedback may be provided on the ATC Intended routing and constraints encountered by an individual flight. An AU may use that information to better select the route yielding a lower cost outcome. For example, as shown in Table 1, a flight has two route choices between a city-pair. Each route choice has an estimated cost based on the plan and the additional ATC-imposed routing and constraints. If optimizing without feedback, the lowest planned cost would be chosen (Route 1 versus Route 2). The feedback allows a more informed choice (Route 2) which lowers the cost. In this case, the benefit of the feedback is the difference in the outcome ($9841- $9681 =$160).

TABLE I. EXAMPLE COST OF ROUTE CHOICE

Case	Route Choice Costs	
	Route 1	***Route 2***
Planned	$ 9660	$ 9671
ATC Intended	$ 9841	$ 9681

Such a feedback mechanism was simulated on the route choice across all flights considered. The distribution of cost savings is shown in Figure 7. Note that 60% of flights get no benefit (not shown in Figure) as the route choice is not affected by the feedback. Despite this, there is an average 0.2% savings in total cost resulting from this feedback. A significant fraction of flights (16%) obtain benefits exceeding 0.5%.

When the CI is set to zero for all flights, the effect on fuel savings alone can be determined. The plots for relative fuel savings looked very similar to the cost savings. On average there was a 0.19% improvement in total fuel consumed when fuel was the optimizing variable.

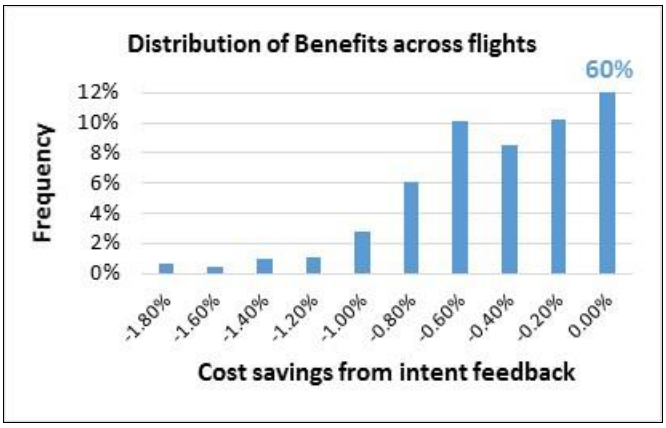

Figure 7. Distribution of cost savings across flights due to feedback of ATC Intended route and constraints (60% get no change)

Figure 8 illustrates that benefits are dependent on the city-pair. Some markets experience very little gain, despite having larger differences between the planned and ATC Intended as shown in Figure 6. There are several factors contributing to these effects such as: number of routes available, cost differences between route choices, and variability in the difference between planned and ATC Intended.

The provision of feedback yields benefits when comparing the cost of the optimal route based on the ATC Intended versus the planned route. However, more advanced AUs may already have acquired route-specific historical knowledge of the impact of ATC Intended routing and constraints. Thus the average effect, for each route, may be applied to the planned route to obtain a better route choice. Returning to the example in Table 1, it may be known that the average impact is an additional 1.4% and 0.8% on routes 1 and 2 respectively. If so, then this historical knowledge would allow the AU to pick Route 2. In this instance there would be no additional benefit to flight-specific feedback.

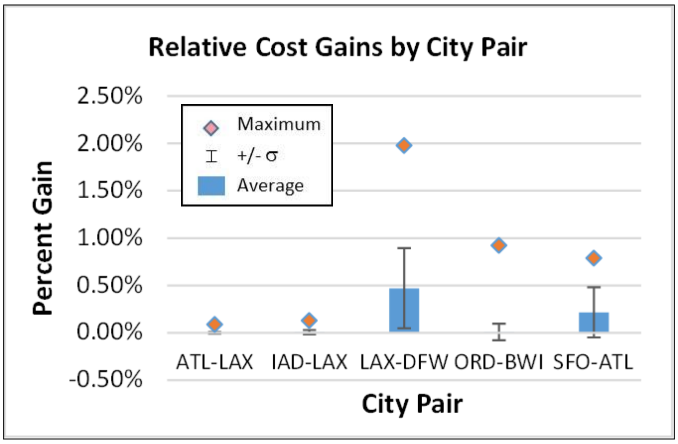

Figure 8. Benefit of feedback by city pair

When route-specific historical knowledge is applied, the average cost reduction from flight-specific feedback is reduced to 0.15% (from 0.2%). Figure 9 illustrates the distribution of benefits across flights and Figure 10 the impact by city pair.

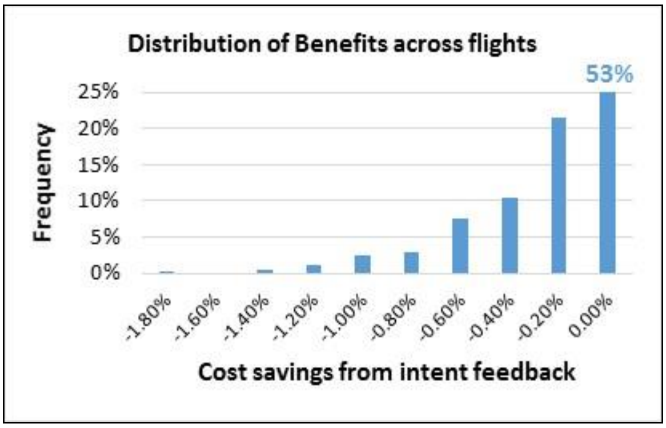

Figure 9. Distribution of cost savings from a base case asuming route-specific historical knowledge across flights due to feedback of ATC Intended route and constraints (53% get no change)

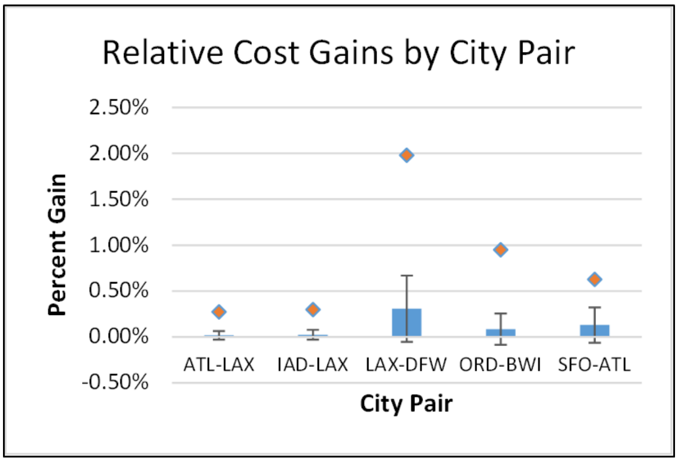

Figure 10. Benefits by city pair when route-specific historical knowledge is applied in the base case

The above illustrates that the application of route-specific historical knowledge reduces the benefit due to flight-specific feedback. However, it does not do this for all cases. As seen in Figure 10, for some city-pairs, flight-specific feedback delivers greater benefits than for the case shown in Figure 8. What is happening is that the impact of ATC Intended routing and constraints is not static. In this case, the use of a mean, without conditioning on additional variables does not help in selecting a better optimum. Figure 11 shows, across a collection of routes, the mean and standard deviation of the cost impact of ATC Intended routes and constraints. The Figure shows that routes sometimes encounter standard deviations comparable in magnitude to their mean.

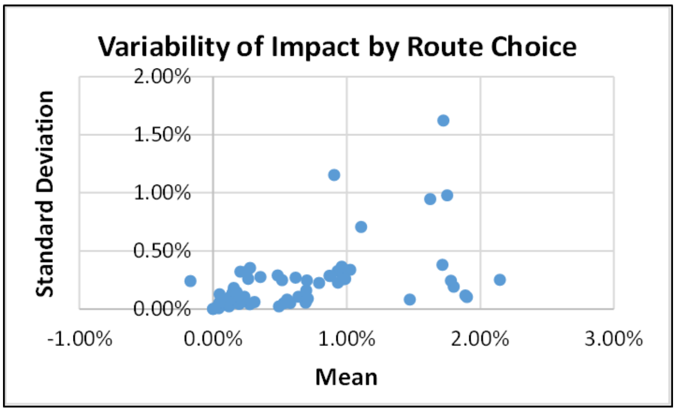

Figure 11. Each dot represents the mean and standard deviation of the relative cost impact of the ATC Intended route and constraints.

Using data from the Bureau of Transportation Statistics [18], major reporting U.S. carriers consumed 50.1 billion dollars of fuel in 2013. A 0.19% improvement in fuel consumed represents a 95 million dollar reduction in fuel consumed. Factoring in the reduction in total costs including fuel and time, the total aircraft operating expenses for reporting U.S. carriers was 89 billion dollars. The magnitude of the gain described herein represents a 178 million dollars per annum cost savings for these carriers from the case without route-specific historical knowledge and 135 million dollars per annum for those where historical knowledge is already applied.

As a validation of the method used to infer cost index, we note that the ratio of total cost to fuel cost across all flights was 1.86 using the inferred cost index. In comparison, the ratio of reported aircraft operating expense to fuel cost was 1.78. This indicates on average that the inferred cost index approach provides an estimate of total cost within 5% given knowledge of flight time and fuel consumption.

B. Results for Ground/Air Delay Optimization

A base case was first simulated representing present-day operations in which no flight is provided with credits for ground delay. In this case, a total of 672 flights were simulated using scheduled demand data from June 1, 2014 across one day. Flights were first modeled deterministically (no EET uncertainty and no departure delay) and verified against a simple spreadsheet queueing model. In this deterministic case, average airborne delay on arrival was computed at 6.2 minutes. The addition of the departure time uncertainty and en route uncertainty results in an average delay of 4.8 minutes when simulated over 100 runs. This computed delay compares to the average 4 minute delay in summer months expressed several years back in a 2008 report [19]. Note that the simulation includes a single averaged capacity value not dependent on fleet mix. The distribution of delays across flights for both cases is illustrated in Figure 12. As expected, delays are reduced when uncertainty is introduced as demand is more spread out with uncertainty.

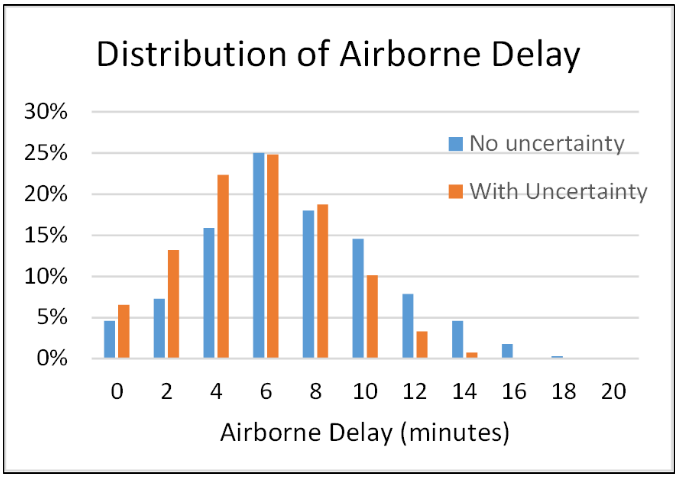

Figure 12. Distribution of simulated airborne delays across flights in the deterministic case

The simulation was run including the process for ground delay allocation described previously. When the ground delay was allocated at the 10th percentile for airborne holding, airborne holding was reduced by 29%; however, total delay (ground and air) increased by 28%. When the cost difference between airborne and ground delay is considered, the net cost of all delay was unchanged between both cases.

An investigation into flight-specific delays revealed why the above approach was not yielding benefits. When a flight takes a ground delay as described, the 10th percentile indicates that there is a 10% chance that the flight will take more ground delay than airborne delay. In this case, the flight will miss the slot it would have obtained without having taken ground delay. Normally there is enough demand in place to take the slot, resulting in no loss in throughput and increase in total delays. However, when the flight having taken too much ground delay occurs at the beginning of a push, there is no demand to take its place. The consequence is significant as every flight behind it experiences additional delay until the next demand gap.

The above situation was confirmed by looking at flight-specific data comparing the case both with and without ground delay allocation. In one case, two flights with one minute each of excess ground delay led to 78 flights having a total of 161 minutes of total additional delay.

Several possibilities were considered to remedy the above situation: 1) only apply the ground delay to flights which are operating during the middle of a push, 2) apply the ground delay to the certain portion of the airborne delay, or 3) consider the probability of system-wide delay versus the single-flight delay. The first approach proved difficult due to the magnitude of departure and flight time errors. The last approach was considered difficult to implement, but may be the subject of future work.

We implemented the second approach by setting the ground delay equal to the minimum airborne delay encountered during the inner Monte Carlo process. This represents the case where a minimum airborne delay is always taken by the flight being considered. By capturing the minimum, we avoid the case where the flight takes too much ground delay.

The average total, airborne and ground delay was computed for each individual flight across multiple Monte Carlo runs. The distribution of individual flights' average airborne and total delay changes is illustrated in Figure 13.

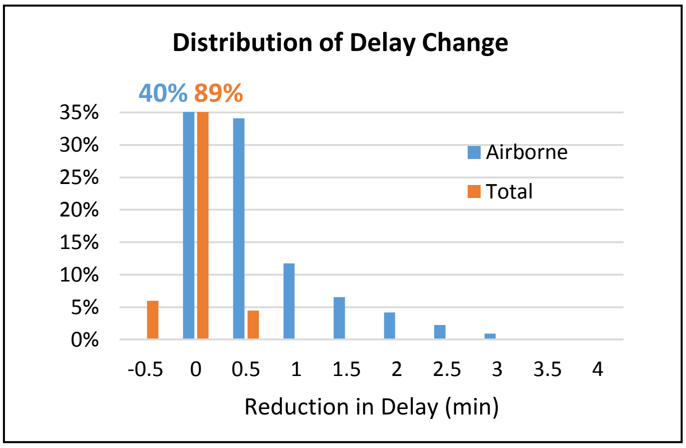

Figure 13. Distribution of changes in delay due to the pre-departure application of ground delays through feedback

Results indicate that airborne delay is reduced by 20 seconds on average. The average sum of ground and airborne delay is slightly increased (11 seconds). Note that a large fraction of flights get no change leading to the small averages. The averaged standard deviation of airborne and ground delay across flights was unchanged at 2.5 minutes. This variation only includes that due to ground and airborne delay and does not include the variation due to EET or departure time uncertainty.

When ground delays are taken by a flight, these are on average 2.6 minutes. When these flights reach their destination, they then receive an average of 4.1 minutes of airborne holding delay. In contrast, when these same flights operate without ground delay due to feedback on minimum expected airborne holding, they experience an average airborne delay of 6.4 minutes. For the case with feedback, the sum of ground and airborne delay is shown, along with the ground delay in Figure 14. For the case without ground delay, airborne only delay is shown. Note that the distribution of total delay taken is approximately the same with and without feedback. However, delays have been taken on the ground for the case with feedback.

The reduction in airborne holding allows flights to save significant amounts of fuel. For example a Boeing 737 holding at 10000 feet would consume approximately 125 kg of additional fuel, and a Boeing 777 would consume 335 kg of additional fuel compared to taking the delay on the ground. With 138 affected flights per day, this leads to a potential savings of 5-13 million kg of fuel per annum by shifting delay to the ground.

The impact of the EET uncertainty was also evaluated to determine the extent to which improved trajectory prediction could improve the allocation of delays further. A reduction in EET uncertainty allows those flights still on the ground to have better situational awareness of the flights in the air when allocating their ground delay. Figure 15 illustrates the trend. While total delay remains almost constant, a 50% reduction in EET uncertainty doubles the ground delay allocation averaged over all flights. This would double the potential fuel savings described previously. A total elimination of EET uncertainty enables a five-fold increase in allocation to ground delay. The remaining airborne delay is required to compensate for ground delay uncertainty.

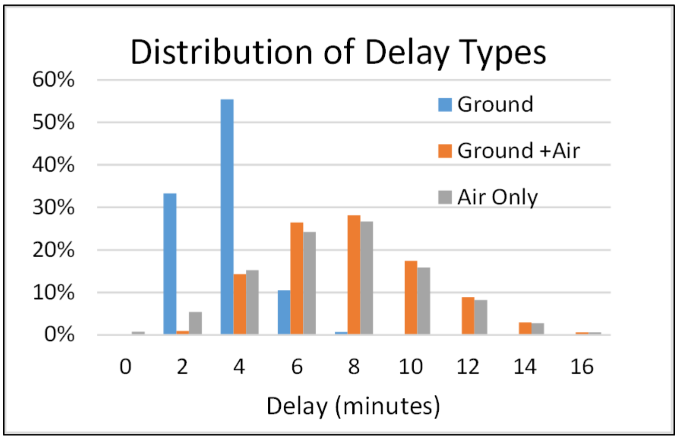

Figure 14. Distribution of delay types for those flights taking ground delay due to feedback on minimum expected airborne delay

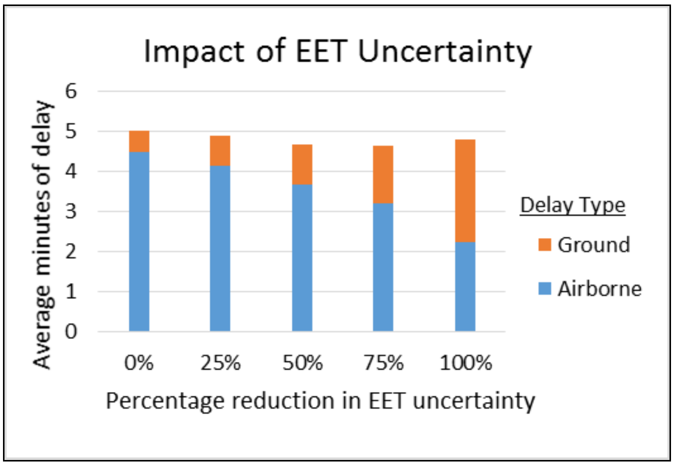

Figure 15. Impact of reduction of EET uncertainty on delay allocation

V. Conclusions

This report evaluated two benefit mechanisms attributable to the exchange of 4DT feedback. The proposed feedback is made possible through changes to flight planning provisions being discussed at ICAO for FF-ICE Step 1.

The first benefit mechanism involves the provision of feedback of applicable re-routing, altitude and speed constraints in a timely manner enabling alternative route choices to be made by the AU. When provided with this feedback, the AU may select the more optimal route resulting in a reduction of 0.2% in the average estimated planned cost of the flight. This gain was dependent on the city pair with some city pairs experiencing near zero change, and one city pair experiencing 0.5% benefits. These small percentage changes can represent hundreds of millions of dollars in cost savings per annum to U.S. carriers.

For advanced AU that may already be taking into consideration the average impact of ATC routing and constraints on planned routes, the average benefit is reduced from 0.2% to 0.15%. This continues to represent a significant cost savings due to flight-specific feedback enabled through implementation of FF-ICE Step 1.

Through the use of multiple route-trajectory pairs enabled through FIXM, flight-specific feedback may be provided on a collection of candidate trajectories. This enables the AU to select the best option in a parallel fashion, versus resorting to multiple serial submissions.

The above analysis investigated the impact of feedback of the re-routes and constraints that are known to automation (e.g., ε in (1)). Further gains might be made by conditioning the distributions of unknown perturbations (e.g., the η in (1)) based on feedback. This is the subject of future work.

The second benefit mechanism explored provided feedback on the distribution of expected airborne delay on arrival. By using this distribution, airspace users can elect to take ground delay on departure in lieu of airborne delays on arrival. A mechanism for crediting these arrivals would have to be in place. This paper did not explore the feasibility of such a mechanism.

The case examined involved a difficult one in which demand and capacity are approximately balanced. Small shifts in demand due to departure and en route uncertainty can cause short-term queueing delays to vary significantly. Aircraft that take too large a ground delay can create a demand gap resulting in delays accumulating against many flights in a push. In this case, selecting a ground delay equal to the lowest expected airborne delay (given all the known information) allows flights which are certain to experience airborne delay to take it as ground delay at a lower cost.

When capacity is degraded, larger delays are anticipated. In such an environment, flights are typically controlled through TFM measures. For those flights that cannot be controlled (e.g., some international arrivals), the feedback mechanism provides a means to incentivize airspace users to voluntarily take a ground delay.

The approach enabled an average of 138 flights per simulated day to shift an average of 2.3 minutes of delay to the ground from airborne delay. On an annualized basis, this represents a pool of 5-13 million kg of fuel per annum for this one airport. Improvements in trajectory prediction accuracy could increase this pool of benefits.

As follow-on work on this benefit mechanism, one should consider allocating ground delay in accordance with a minimum system-level cost. This would allow the distribution of total cost to consider the impact one flight has on the push, versus simply on itself.

The example of London Heathrow was not picked to demonstrate any specific benefits at that airport as certain flow management measures are already in place to mitigate airborne holding. However, other international airports experiencing similar airborne holding may not presently have the ability to influence departures. Through FF-ICE Step 1, and pre-departure feedback, a dialogue between airspace users and the ATM service providers allows the airspace users to take steps to mitigate the holding.

Through the exchange of pre-departure feedback, the provisions being developed have the potential to improve flight operations. Feedback during planning can improve the selection amongst available choices. When operating more tactically, feedback on the expected impact of the most up-to-date information can enable tactical choices such as taking ground delay in lieu of airborne delay.

ACKNOWLEDGMENT

The authors would like to acknowledge contributions from Ruby Sharma (MITRE), Harriet Roberts and Sean Megeath (both from Vencore). Discussions with Philippe Leplae and Kim Breivik of EUROCONTROL under Coordination Plan 3.1 have yielded significant progress towards the harmonization of the use of the 4DT in FF-ICE Step 1.

REFERENCES

[1] ICAO Secretary, Memorandum 1 to the Flight Plan Study Group members, June 2004

[2] International Civil Aviation Organization, Proedures for Air Navigation Services – Air Traffic Management, ICAO Doc. 4444, 15th Edition, Incorporating Amendments 1-5, 2007

[3] International Civil Aviation Organization, Manual on Flight and Flow Information for a Collaborative Environment, ICAO Doc. 9965, 2011

[4] International Civil Aviation Organization, Twelfth Air Navigation Conference, Working Paper 158, Draft report on Agenda Item 3, November 2012

[5] International Civil Aviation Organization, 2013-2028 Global Air Navigation Plan, ICAO Doc.9750-AN/963, 4th Edition, 2013

[6] Flight Information Exchange, web site, http://www.fixm.aero

[7] K. Lindsay, D. Greenbaum, C. Wanke, "Impact of Pre-Departure Uncertainty on Collaborative Routing Coordination Tools (CRCT) Prediction Performance," AIAA 3rd ATIO, Denver, CO, 2003

[8] J. Krozel, D. Rosman, S. Grabbe, "Analysis of En Route Sector Demand Error Sources," AIAA-2000-5016, AIAA GN&C Conference, Monterey, CA, 2002

[9] C. Wanke, S. Mulgund, D. Greenbaum, L. Song, "Modeling Traffic Prediction Uncertainty for Traffic Management Decision Support," AIAA-2004-5230, AIAA GN&C Conference, Providence, RI, 2004

[10] R.A. Coppenbarger, "Real-Time Data Link of Aircraft Parameters to the Center-TRACON Automation System (CTAS)," 4th USA/Europe ATM R&D Seminar, Santa Fe, NM, 2001

[11] D.H. Williams, S.M. Green, "Flight Evaluation of Center-TRACON Automation System Trajectory Prediction Process," NASA Tech Report, TP-1998-208439, July 1998

[12] C. Wanke, "Using Air-Ground Data Link to Improve Air Traffic Management Decision Support Tool Performance," USA/Europe ATM R&D Seminar, Saclay, France, 1997

[13] W.C. Arthur, M.P. McLaughlin, "User Request Evaluation Tool (URET) Interfacility Conflict Probe Performance Assessment," 2nd USA/Europe ATM R&D Seminar, Orlando, FL, 1998

[14] S. Mondoloni, "Improved Trajectory Information for the Future Flight Planning Environment," 10th USA/Europe ATM R&D Seminar, Chicago, IL, 2013

[15] T. Gaydos, B. Liao, E. Smith, L. Wang, "Analysis of Conflict Detection Performance for Trajectory-Based Descent Operations," 10TH usa/Europe ATM R&D Seminar, Chicago, IL, 2013

[16] J. Rakas, B. B. Niu, J. Tom, C. Santiago, M. Seelhorst, "Analysis of Air Traffic Control Command Entries and the Impact on Decision Support Tool Performance," *Air Traffic Control Quarterly*, Vol. 22, No. 1, 2014

[17] M. K. DeJonge, W.H. Syblon, "Application of Cost Index to Fleet Hub Operation," American Control Conference Proceedings, pp.179-183, San Diego, CA, 1984

[18] Department of Transportation, Bureau of Transportation Statistics website http://aps.bts.gov

[19] SH&E Consulting, UK CAA Runway Resiliency Study – Final Report, December 2008

AUTHOR BIOGRAPHY

Stéphane L. Mondoloni received an S.B., S.M., and Ph.D. in 1987, 1989, and 1993 respectively from the Massachusetts Institute of Technology in Cambridge, MA, USA. All degrees were in aeronautical and astronautical engineering. Dr. Mondoloni is an Associate Fellow of the AIAA, presently serves as an advisor to the US member for an ICAO Panel developing future flight planning provisions, and works as a Senior Principal Engineer at the MITRE Corporation in McLean, VA, USA. His current research interests include ATM performance evaluation, trajectory-based operations and flight planning.

Sheng S. Liu received a B.S. in Computer Science from State University of New York at Binghamton and a M.S. in Computer Science with Bioinformatics concentration from Johns Hopkins University. He works as a Senior Multi-Discipline Engineer at the MITRE Corporation at Mclean Virginia, USA. His research interests include ATM research computing capability, modeling and simulation, and future flight planning and ATFM concepts.

Daniel Kirk received his Ph.D. from the University of Michigan in 1985. He joined MITRE/CAASD in 1988, where he is currently a Senior Principal Engineer. His work at MITRE has focused on the development of advanced air traffic decision support capabilities, including the User Request Evaluation Tool (URET) conflict probe. His current work is focused on the development and assessment of trajectory modeling enhancements, flight information exchange, and integration of decision support capabilities at the sector.