

Benefits Analysis of a Departure Management Prototype for the New York Area

*James DeArmon, Norma Taber, Hilton Bateman,
Lixia Song, Tudor Masek*
The MITRE Corporation
Center for Advanced Aviation System Development
McLean, Virginia, USA

Daniel Gilani
Federal Aviation Administration
Washington, D. C., USA

Integrated Departure Route Planning (IDRP) is a decision support tool being developed and prototyped by MITRE/CAASD to explore new concepts and capabilities for departure management. IDRP provides demand estimates for departure fixes and routes in terminal airspace, including identification of specific flights impacted by capacity constraints, their route information, and accurate estimates of their expected take-off times.

In general, IDRP benefits accrue when there is contention for departure resources (runways, fixes, routes, sectors), plus the feasibility of off-loading or otherwise balancing demand as a means of mitigating delay. This scenario is common in the New York area, where the prototype has been installed since 2009. In 2011 and 2012, field evaluations were conducted at tower, terminal, and center facilities. These evaluations allowed the capture of “use cases”—instances of essential applications of the tool. These use cases were later examined via offline replay, and led to benefits analyses in which a queuing model was employed to compare scenarios with and without IDRP. The modeling suggests significant benefits are attributable to IDRP.

I. INTRODUCTION

Air traffic departing from New York area airports in the U.S. experience some of the worst ground-side and airspace congestion in the National Airspace System (NAS) [1]. Departure traffic management has the challenge of balancing departure demand with available capacity, and seeks to ameliorate this congestion and expedite traffic movement. There is a need for improved automation tools that provide integrated data sources (such as traffic demand, weather impacts, and airspace availability) so that departure traffic management can effectively execute its tasks.

Integrated Departure Route Planning (IDRP) is a prototype in use at air traffic control facilities as well as some airline operations centers for departures from New York airports. The prototype was first put into use in 2009. This paper describes field evaluations in 2011 and 2012, and provides benefits analyses of observed use cases.

II. BACKGROUND

The IDRP decision support capabilities are being developed and prototyped by The MITRE Corporation’s Center for

Advanced Aviation System Development (MITRE/CAASD) in collaboration with Massachusetts Institute of Technology’s Lincoln Laboratory (MIT/LL). IDRP augments the capabilities of MIT/LL’s Route Availability Planning Tool (RAPT) [2], with route and fix demand information. IDRP combines state-of-the-art weather forecasts with operational flight data in an effort to assist Federal Aviation Administration (FAA) air traffic managers and commercial flight operators in making proactive Traffic Flow Management (TFM) decisions, both during severe weather and during clear weather conditions when traffic demands are reaching or exceeding the capacity of NAS resources. Fig. 1 shows the IDRP graphical user interface.

The IDRP prototype was developed for the purpose of conducting field evaluations. The vision is for the IDRP capabilities to be incorporated into the Collaborative Air Traffic Management – Technologies (CATM-T) program [3], as part of a “mid-term” (2017-2020) functional enhancement package, thereby helping to fulfill goals of the Next Generation Air Transportation System (NextGen) [4].

IDRP’s coverage of NextGen needs can be categorized in three key ways. First, IDRP capabilities provide near-term predictions of the impact of weather on flight routing. They also provide demand estimates for departure fixes and defined flight routes within the New York airspace. While legacy systems provide accurate estimates of weather severity in this airspace, the IDRP capabilities identify the specific flights that will be impacted by those constraints and provide more accurate departure time estimates for these flights.

Second, IDRP disseminates information indicating the impact of congestion and weather on flight routes and suggests potential flight-specific trajectory changes to FAA and flight operator facilities. This information is expected to improve the predictability of reroutes, allowing for proactive decision-making and efficient re-planning.

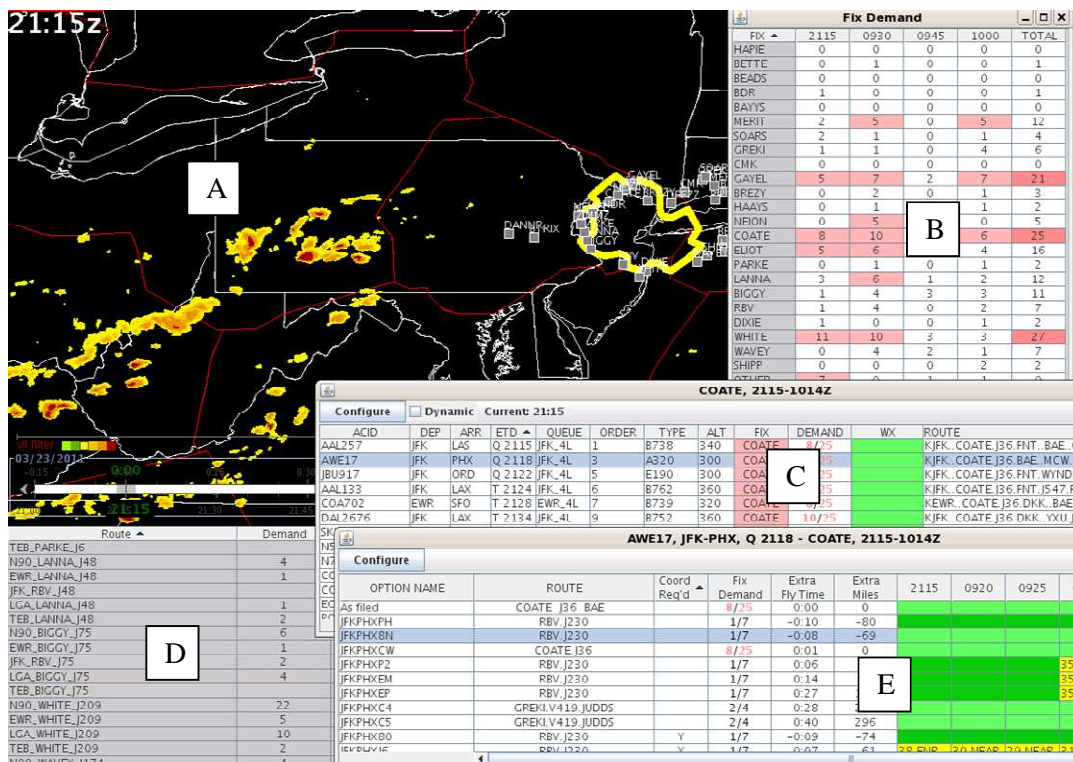


Figure 1. IDR P Graphical User Interface. Five windows are displayed: A. Integrated Traffic and Weather Map, B. Fix Demand Table, C. Flight List, D. Route Impact and Demand Table, E. Reroute Options List

Finally, by providing common situational awareness across the FAA and flight operator facilities of NAS conditions and available reroutes, IDR P supports timely, effective, and informed decision making. Because personnel in the outfitted facilities have access to the same information, IDR P allows decision makers to work together to exchange relevant information quickly and to support the right person making the right decision at the right time.

Research on and evaluation of the IDR P capabilities are being performed in an incremental fashion. Initially, IDR P capabilities were implemented as enhancements to the RAPT system in the New York area. Phase 1 of IDR P included enhancements to RAPT which provide demand, capacity, and alert information on departure fixes and routes. Phase 2 includes flight-specific reroute recommendations to the user. The IDR P Phase 2 capabilities were in use during the 2011 and 2012 field evaluations.

III. FIELD EVALUATIONS IN 2011, 2012 AND OBSERVED USE CASES

During the summer of 2011, over 2000 observations spanning 74.5 hours were recorded at multiple New York facilities. The facilities observed included New York Center (ZNY) and New York Terminal Radar Approach Control (TRACON), as well as airport towers and flight operations centers. These observations were grouped into various use cases and the participants at ZNY and New York TRACON (N90) were asked to confirm that these observation groupings

represented typical traffic management events. Prior to the 2012 evaluations the ZNY and N90 participants were asked to provide a list of the five most common events observed in 2011 that were likely to benefit from IDR P's functions and capabilities. These events became the focus of the evaluations during 2012, which also identified additional events during 63 hours of observations. From the full set of 2011 and 2012 observations, three events that illustrate the potential benefits of IDR P were chosen for further analysis, as detailed in this paper. These three events are examples of the following use cases:

- Offloading demand from a saturated fix: During certain periods of the day, New York area departure fixes experience condensed demand from a significant number of flights requesting to depart multiple airports at the same time over the same departure fix. If unmonitored, this situation leads to a high volume of traffic in the N90 departure area and results in multiple airport departure stops, or stringent miles-in-trail (MIT) or minutes-in-trail (MINIT) restrictions. Traffic managers apply tactical routing of aircraft away from the congested fixes to prevent condensed departure demand from occurring.
- Balancing departure runway loads: Another technique for managing departure fix congestion is to apply additional spacing between flights over a congested fix when the congestion results from aircraft departing from the same airport. Assignment of these (MINIT)

restrictions is usually done uniformly across the major airports, without a detailed understanding of the need for such restrictions and how those restrictions affect airport surface operations.

- Combining fix offloads with diverging-heading departures: The normal departure configuration for LaGuardia Airport (LGA) allows for only one departure between each pair of arrivals. Often this results in departure delays when the rate of aircraft becoming ready to depart exceeds the departure capacity. This can result in surface gridlock, which must be alleviated by putting arrival traffic into holding patterns. A modified departure configuration that can help alleviate the departure delays is possible when two sequential departures are going over fixes that are sufficiently separated that the departing flights will be on diverging courses. Rerouting a few flights to provide this alternating departure flow allows two departures between each arrival pair, thereby reducing departure delays.

IV. ESTIMATING BENEFITS

An important activity in deploying a new capability is an assessment of operational benefits. Credible benefits analyses assure developers that their efforts are worthwhile, and help justify continued program funding. A challenge with estimating benefits of IDRP is the development of baseline vs. treatment cases for comparison. From the field evaluations, observed use cases were found in which IDRP was in use, representing a treatment case. A real-world baseline case for operations without IDRP does not and cannot exist; i.e., it is impossible to find examples where all conditions of the real-world treatment case were in effect *except* for the IDRP usage. We therefore resort to modeling, an abstraction of the real world, wherein it is easy to compare identical cases with IDRP and without IDRP.

A. Application of a Simple Queuing Model

In our analysis, estimation of delay associated with resource over-subscription was accomplished via a simple queuing model. For some number of hours, capacity and demand are supplied as input, and compared. For each hour, if demand is less-than-or-equal-to capacity, then no delay occurs. But if demand exceeds capacity, then unsatisfied demand “spills over” into the next hour. Not only do the spill-over flights accrue delay, but, consistent with a first-come, first-served discipline, flights in the next hour also suffer delay, as they get pushed back by the spill-over flights.

Both demand and capacity are assumed to be uniformly distributed in an hour. For example, if capacity is 15 per hour, then there is a 4-minute ($=60/15$) service time per flight. This modeling approach was promulgated in [5] and validated in an air traffic management application by MIT/LL [6]. Although this model is abstracted and simple, it is considered a reasonable means of representing a “noisy” system. The results are almost certainly understated, since other flights not represented in the model experience delay in the real-world,

due to their being in line behind flights with take-off restrictions.

Three observed cases are presented below, with their accompanying hypothesized No-IDRP case. The difference in delay between No IDRP and With IDRP represents a benefit of IDRP. Some further assumptions and calculations are employed to monetize and annualize these results.

The monetized savings of avoiding delay is defined as follows:

- A flight minute of ground delay is valued at \$36, which represents airline direct operating costs (ADOC) [7].
- Passenger value of time (PVT) is valued at \$32 per flight minute (assuming nominal passenger load factors) [8]

We examine the queuing results for an operational example of each of the three use cases described earlier.

B. Case 1: Offload from Saturated Fix

On 19 July 2011 at 2000 Greenwich Mean Time (GMT), traffic managers used IDRP to determine that the ELIOT departure fix was projected to have excess demand, in light of a 2-hour 20 MIT restriction. A nearby fix, COATE, was shown to have sufficient available capacity to accommodate flights offloaded from ELIOT. The Supervisory Traffic Management Coordinator (STMC) offloaded 4 flights in the 2000 hour and 1 flight in the 2100 hour. Tables I and II show the demand and capacity values by hour for 3 hours. In Table II, it is assumed that the offloads were not performed.

TABLE I. CASE 1, WITH IDRP: DEMAND (AFTER OFFLOADS) AND CAPACITY FOR 3 HOURS

Hour (GMT)	2000	2100	2200
Demand	23	18	18
Capacity (20 MIT*)	15	15	24**

* A nominal flying speed when crossing ELIOT is 300 knots – 5 nautical miles (nm)/minute. That means it takes 4 minutes to go to 20 NM. So with a 20 MIT restriction, flights are 4 minutes in trail, achieving an hourly rate of 15 (+60/4).

** The full capacity of ELIOT is assumed to be 24/hour per consultation with subject matter experts.

TABLE II. CASE 1, NO IDRP: ASSUME OFFLOADS NOT PERFORMED

Hour (GMT)	2000	2100	2200
Demand	27	19	18
Capacity (at 20 MIT)	15	15	24

The queuing model results: 790 flight minutes of delay for No IDRP, 347 for With IDRP, for a savings of 443 flight minutes.

There was a possible cost, however, in terms of additional air miles when routed over the alternate fix. An analysis was performed, examining flight paths over ELIOT vs. over

COATE, for matched origin/destination pairs. For various departure flight pairs, the flight distance was compared. Depending on the origin and destination, there were small differences in flight distance, over ELIOT vs. over COATE, but neither routing was clearly shorter. No adjustment to the 443 flight minutes of savings was justified.

Another consideration regarding accuracy of these benefits is the definition of the No-IDRP case. It could likely be the case that, even without IDRP, some offloads would have been performed. It is difficult, however, to conjecture how many. It is hoped that conservative assumptions in other parts of this analysis will balance out this point.

Valuing the 443 flight minutes at \$36/minute, this is a savings in ADOC, for this one situation, of \$15,948. In terms of PVT, this is a savings of \$14,176.

C. Case 2: Departure Runway Load Balancing

On 17 June 2011 at 1915 GMT, severe en route weather was impacting ZNY and surrounding centers. A Severe Weather Avoidance Plan (SWAP) was implemented with associated traffic flow management actions. Some of these actions closed multiple routes, and many flights were rerouted over ELIOT. Since ELIOT was unaffected by weather, it could run at full capacity. A typical traffic management solution in this situation is to apply 5 MINIT separately to all airports with departure flights over ELIOT, in this case LGA and Newark Liberty International Airport (EWR). However, it was observed that the STMC used IDRP to examine the relative queue sizes at EWR and LGA for ELIOT departures, noted a severe imbalance, and implemented 7 MINIT for EWR departures and allowed LGA departures to “free flow,” i.e., depart over ELIOT unrestricted. Table III shows the demand/capacity situation for LGA Runway 13 for the subject time period. By contrast Table IV reflects the typical traffic management initiative (TMI) of 5 MINIT for LGA departures over ELIOT. (It happened that EWR had so little ELIOT demand that the 7 MINIT restriction added no additional delay, and it was therefore not necessary to model for this analysis.)

TABLE III. CASE 2, WITH IDRP: LGA RUNWAY 13 DEPARTURES IN LIGHT OF “FREE FLOWING”

Hour (GMT)	1915	2015
Demand	29	26
Runway Capacity	30*	30

* A nominal departure rate of 30 per hour, i.e., 2-minute spacing is assumed.

TABLE IV. CASE 2, NO IDRP: NOMINAL FLOW MANAGEMENT ACTION OF 5 MINUTE-IN-TRAIL FOR LGA DEPARTURES OVER ELIOT

Hour (GMT)	1915	2015
Demand	29	26
(Effective) Runway Capacity	17*	30

* This capacity value was computed manually: the 9 back-to-back pairs over ELIOT each get an additional 3 minutes spacing. Nine pairs times 3 additional minutes = 27 additional minutes of spacing,

nearly halving the departure runway capacity.

The queuing model results are shown in Table V.

TABLE V. CASE 2 RESULTS

Savings (No IDRP minus With IDRP)	ADOC Valuation	PVT Valuation
228 (=228 – 0) flight minutes	\$8,208	\$7,296

D. Case 3: Offload to Fix, Plus Diverging Departures

At approximately 1345 GMT on 4 June 2012, Departure Sequencing Program (DSP) delays at LGA were shown to be exceeding 30 minutes and IDRP showed 25 aircraft in the LGA departure queue. The runway configuration at LGA required departures on Runway 13 and arrivals on Runway 04, typically a “one in, one out” operation. The Tactical Route Coordinator (TRC) at N90 examined IDRP’s LGA flight list using the Flight List feature, and recognized that the departure demand over the BIGGY departure was heavy, whereas departure demand over the RBV fix was light. The TRC realized that moving some departures from over BIGGY to over RBV would enable the tower staff to use both the TNNIS and CONEY climbs, which are Standard Instrument Departure (SID) procedures. This would greatly increase the potential departure throughput of LGA by enabling the tower to launch two departures, on diverging headings, between successive arrivals when the departure flight sequencing allowed this.

The delay associated with this departure runway congestion was estimated with the simple queuing model. To set up the model, consider in Table VI the observed LGA departure counts for that date (as taken from FAA CountOps data [9])

TABLE VI. OBSERVED LGA DEPARTURE COUNTS FOR 12 JUNE 2012

Hour	1200	1300	1400	1500	1600	1700
Count	34	36	35	40	36	36

Since there was excess departure demand during these hours, as evidenced by long runway queues, these counts may be considered as departure runway capacities for modeling purposes, i.e., maximum rates possible, given the conditions. Note the capacity increase during the 1500 hour—the diverging heading operations reduced the inter-departure timing. A departure rate of 40 in an hour is not unheard of for LGA, depending on runway configuration, fleet mix, staffing, etc. However for the case at hand, the 40 rate should be considered as an increased hourly departure rate of about 4 or 5 flights, in light of the 35 and 36 rates in the adjacent hours.

A second input needed for the model is departure demand per hour. To estimate these values, two prior weeks (weekdays only) of hourly departure demand (as represented by proposed departure times) were examined. The 22 May 2012 counts were closest to the average counts for the two weeks

considered. The highest demand in the two weeks was on the subject date of 4 June. Demand counts for these two dates are given in Table VII.

TABLE VII. LGA DEPARTURE DEMAND COUNTS PER HOUR

Hour	1200	1300	1400	1500	1600	1700
22 May	47	35	34	37	36	35
4 June	50	36	37	39	33	39

For modeling purposes, pursuant to more generalized results, we will use the 22 May 2012 counts.

The scenario for With IDRP is shown in Table VIII. The increased rate in the 1500 hour is in bold font. The scenario for No IDRP is shown in Table IX. Note that for the 1500 hour, capacity is assumed to be reduced by 4.

TABLE VIII. CASE 3, WITH IDRP: OFFLOAD LGA DEPARTURES TO RBV, WITH DIVERGING DEPARTURES

Hour	1200	1300	1400	1500	1600	1700
Demand	47	35	34	37	36	35
Capacity	34	36	35	40	36	36

TABLE IX. CASE 3, NO IDRP: NO OFFLOADS

Hour	1200	1300	1400	1500	1600	1700
Demand	47	35	34	37	36	35
Capacity	34	36	35	36	36	36

The queuing model results are shown in Table X.

TABLE X. CASE 3 RESULTS

Savings (No IDRP minus With IDRP) flight minutes	ADOC Valuation	PVT Valuation
214 (=765 – 551)	\$7,704	\$6,848

V. COMPARISON WITH DISCRETE-EVENT SIMULATION

An alternate modeling method, being developed to support “what-if” modeling of alternatives within the IDRP prototype, was also applied to one of the use cases presented. This approach uses a discrete event simulation to account for the individual flights in a given departure scenario and to separate the flights appropriately based on nominal runway spacing requirements, as well as any additional airspace constraints such as TMIs. By modeling individual flights queuing on the ground, this approach captures the secondary impacts of flights filed over congested departure resources blocking other flights behind them waiting to depart.

To properly capture the operational resources available at various airports, the discrete-event simulation can model multiple departure runways per airport and can also model multiple “feeder” taxi queues for each of those runways. As new flights approach their estimated departure times, they are assigned a departure runway and subsequently join a specific taxi queue. The default model logic is that, when there are multiple departure runways available, flights are more likely to be assigned to the runway with the longest current departure queue. This in effect simulates a “primary” departure runway, such as at John F. Kennedy International Airport (JFK). However, when the assigned runway has multiple taxi queues, flights are more likely to be assigned the shortest such queue, thus simulating taxiway load balancing. More complicated operations—such as assigning all flights filed over some set of departure fixes to a single runway or dedicating one of the available taxi queues solely to impacted flights so that the other can “free-flow”—could be easily captured in the simulation; however, defining all of the airport-specific surface operations only adds to the complexity of the model and increases the number of model parameters to set and validate. The initial comparisons presented here used the default queue assignment logic.

The simulation models various levels of dependencies between departing flights: 1) within queue, 2) across taxi queues for a single runway, 3) across runways for a single airport, as well as 4) across airports. Various rules for time separation of departures and the impact of TMIs were formulated to represent these dependencies.

Finally, all flights are processed on a first-come-first-served basis. When two or more flights have the same estimated departure time, the flight with the longest departure delay goes first. After each departure, all subsequent flights’ estimated departure times are updated to reflect the minimum runway (and possibly TMI) separation requirements.

A. Results

The discrete-event simulation was used to analyze the second use case presented above: departure runway load balancing at LGA on 17 June 2011. The following nominal runway configuration was used: 2 active departure runways at JFK, 1 at all the other N90 airports; 2 taxi queues per runway at JFK, EWR, and LGA; 1 queue at all the other airports. The nominal runway spacing requirements were set to 60 seconds for any two flights and 90 seconds for consecutive flights over the same fix. The set of TMIs active between 1915 and 2115 GMT (taken from the National Traffic Management Log archive [10]) were also loaded and simulated. Departure demand data were taken from the operational data archive. As a final step in the set-up, in order to compare directly to the estimate provided by the simple queuing model, only the Runway 13 departures were simulated at LGA.

In the With IDRP case of no LGA restriction over ELIOT and 7 MINIT from EWR over ELIOT, this simulation approach yielded a total delay across the New York airports of 4,814 minutes. In the hypothetical “No IDRP” case with 5 MINIT over ELIOT for both LGA and EWR (separately), the discrete-event simulation estimated a total of 5,281 minutes of delay.

Therefore, the estimated benefit of load balancing in this case is 467 minutes, about twice the result of 228 minutes savings per the simple queuing model. This simulation exercise may be considered an approximate cross-validation with the simple queuing model. The simulation estimates a higher benefit, compared to the simple queuing model, because it includes flights bound for all the other departure fixes in addition to ELIOT. At EWR and LGA in particular, these other flights experience secondary delay impacts due to their waiting in queue behind the flights bound for ELIOT. Subject matter experts have reviewed both results and deemed them reasonable.

VI. CONCLUSION

This paper presented three use cases that exemplify the application of IDRP and the estimated resultant benefits obtained from a simple queuing model. For one of the use cases, a cross-validation was performed by constructing and exercising a higher-fidelity discrete event simulation model. An expected outcome was the simulation model estimating a greater benefit, compared to the queuing model.

The modeling demonstrates that even moderate intervention by traffic flow managers in a selective and focused manner can yield important delay and cost savings to air carriers. The use cases highlighted here were common situations, occurring every day or at least several times per week. On an annual basis, this comes to hundreds of times for each use case, and therefore would represent significant annual benefit. For example, estimating a single use per business day in New York (260 days per year), with an average savings per day of approximately \$15,000, results in an annual savings of \$3.9 Million in New York. If the benefits were similar at, for example, 4 other NAS metroplexes, a net savings per year of about \$20 Million would result.

REFERENCES

[1] GAO, 2010, "Setting On-Time Performance Targets at Congested Airports Could Help Focus FAA's Actions" [online] <http://www.gao.gov/new.items/d10542.pdf>

[2] MIT/Lincoln Laboratory, "Route Availability Planning Tool" [online] <http://www.ll.mit.edu/mission/aviation/wxatmintegration/rapt.html>

[3] FAA, 2012, "Collaborative Air Traffic Management Technologies" [online] <http://www.faa.gov/nextgen/implementation/programs/catmt/>

[4] Joint Planning and Development Office, 2010, Concept of Operations for the Next Generation Air Transportation System, Version 3.2. Washington, D.C.

[5] G.F. Newell, 1971, Applications of Queuing Theory, Chapman Hall, London.

[6] Allan, S. S., and J. E. Evans, 2005, "Operational Benefits of the Integrated Terminal Weather System (ITWS) at Atlanta," Project Report ATC-320, MIT Lincoln Laboratory, Lexington, MA.

[7] FAA Investment Planning and Analysis, Economic Information for IA. [online] <http://www.ipa.faa.gov/index.cfm>

[8] FAA, 2008, "Benefits Basis of Estimate for Collaborative Air Traffic Management Technologies (CATM-T)," Work Package 2, Department of Transportation, Washington, DC.

[9] FAA, 2011, The Operations Network Reporting Facility [online] http://aspmhelp.faa.gov/index.php/OPSNET_Delays:_Standard_Report

[10] FAA, 2011, National Traffic Management Log [online] <http://hf.tc.faa.gov/projects/ntml.htm>

NOTICE

This work was produced for the U.S. Government under Contract DTFAWA-10-C-00080 and is subject to Federal Aviation Administration Acquisition Management System Clause 3.5-13, Rights In Data-General, Alt. III and Alt. IV (Oct. 1996).

The contents of this document reflect the views of the author and The MITRE Corporation and do not necessarily reflect the views of the Federal Aviation Administration (FAA) or the Department of Transportation (DOT). Neither the FAA nor the DOT makes any warranty or guarantee, expressed or implied, concerning the content or accuracy of these views.

© 2013 The MITRE Corporation. All Rights Reserved.

Approved for Public Release; Distribution Unlimited.
Case Number: 13-0046.