

UC Berkeley

Research Reports

Title

Assessing URET Benefits for Airspace Users: A Quasi-Experimental Approach

Permalink

<https://escholarship.org/uc/item/99h0s0rt>

Authors

Bolic, Tatjana
Hansen, Mark

Publication Date

2002-04-01

Institute of Transportation Studies
University of California at Berkeley

**Assessing URET Benefits for Airspace Users:
A Quasi-Experimental Approach**

Tatjana Bolic and Mark Hansen

RESEARCH REPORT
UCB-ITS-RR-2002-3

NATIONAL CENTER OF EXCELLENCE FOR
NEXTOR
AVIATION OPERATIONS RESEARCH

NEXTOR Research Report 02-1

April 2002
ISSN 0192 4095

Abstract

Air traffic control organizations around the world are trying to develop automation tools to help controllers manage increasing workload and to enable user preferred routes. This paper focuses on such a tool: User Request Evaluation Tool (URET), which is a decision-support tool for en-route controllers. URET is a prototype of an automated conflict probe. Based on flight plans and actual radar tracks, the URET system models aircraft trajectories and predicts possible conflicts. It also enables controllers to check clearances for conflicts prior to their issuance. This tool is intended as a strategic decision-support tool for the D-side controller.

When implemented in the Air traffic control system, tools like URET can affect many aspects of system performance, from controller workload, to safety, to the quality of service provided to users. The purpose of this paper is to determine the impact of URET on users--more specifically on flight times experienced by users. To do so, we employ statistical methods that compare changes in flight times before and after URET implementation for two sets of flights. One set includes flights that traverse the airspace where URET is implemented; the other includes flights that do not use that airspace.

Results suggest that URET reduces flight times for flights using the URET airspace. Departure delay, rather than airborne time, is the flight time component that is most strongly affected, with a decrease of 0.5-3 minutes per flight depending on the analysis approach and time period analyzed. Airborne time reductions are in the range of 0.2-0.5 minutes per flight under one analysis approach, and statistically insignificant under the other approach. These results imply that much of the benefit from URET is non-local and derives from mechanisms other than more direct routing through URET airspace, which has been the focus of most earlier benefits studies.

Table of Contents:

Abstract	I
Table of Contents:	II
List of Figures:	II
List of Tables:.....	III
1. Introduction	1
2. URET Functionality and Deployment.....	2
2.1 URET Functionality	3
2.2.1 URET Build 1.....	3
2.2.2 URET Build 2.....	7
2.3 URET Development.....	8
3. Assessing URET—A Literature Review.....	12
4. Methodology Overview.....	18
5. Average Flight Time Analysis	19
5.1 Methodology	19
5.2 Results	23
6. Individual Flight Time Analysis	31
6.1 Individual Flight Times Model	31
6.2 Model Estimates Discussion	33
6.3 URET Influence on Flight Times.....	38
7. Caveats	46
8. Conclusions	51
References:	52
List of Abbreviations:.....	53

List of Figures:

Figure 1. URET Sites During FFP1 (Source: FFP1 Performance Metrics Team Report, December 2000)	4
Figure 2. URET Deployment.	9
Figure 3. En-Route Distance Trend, ZME (Source: FFP1 Metrics Report, June 2001) ..	15
Figure 4. En-Route Distance Trend, ZID (Source: FFP1 Metrics Report, June 2001)	16
Figure 5. URET and Non-URET City-pairs Generation.....	20
Figure 6. Metrics Calculation Flow-chart.	22
Figure 7. Average Airborne Improvement.	27
Figure 8. Average Flight Time Improvement.	29
Figure 9. Piece-Wise Linear Distance-Time Function.....	36
Figure 10. Influence of Difference in Longitude on East and Westbound Flights.	37
Figure 11. Distribution of Average Daily Delays.	41
Figure 12. Radar Weather Observations for Days 2/13/2000 and 2/3/1999. (Source: http://www4.ncdc.noaa.gov/cgi-win/wwcgi.dll?wwnextrad~images2).....	48
Figure 13. AATC for URET and Non-URET City-pairs.	49

List of Tables:

Table 1. URET Functional Capabilities	5
Table 2. Airports Used in the URET City-pairs and Their Pairs Building Frequency (PBF)	24
Table 3. Airports Used for Non-URET City-pairs and Their Pairs Building Frequency (PBF)	25
Table 4. AATC Metric Values for URET and Non-URET City-pairs.....	26
Table 5. AFTC Metric Values for URET and Non-URET City-pairs	28
Table 6. Individual Flight Time Coefficients for June.....	34
Table 7. Airport Fixed Effects Coefficients for June	35
Table 8. Regression Coefficients of Individual Flight Times Model and Its Components (minutes).....	39
Table 9. Regression Coefficients of Individual Flight Times Model and Its Components, Group 1.....	42
Table 10. Regression Coefficients of Individual Flight Times Model and Its Components, Group 2.....	43
Table 11. Regression Coefficients of Individual Flight Times Model and Its Components, Group 3.....	44
Table 12. Hourly Throughput for ZID and ZME.....	47

1. Introduction

Air traffic control operations today are highly structured and restrictive, requiring aircraft to fly along predefined airways, often at proscribed altitudes. This structured and restrictive system helps air traffic controllers manage their workload while assuring that required separations are maintained. At the same time, however, airspace users desire a less structured, more flexible system with greater freedom to fly direct or wind-optimal routes, and fuel-efficient flight profiles. Such flexibility, if it can be attained without compromising safety, could result in considerable savings in fuel and time.

In order to achieve more flexibility in Air traffic control, the Federal Aviation Administration has deployed several decision support tools at selected sites as part of the Free Flight Phase One (FFP1) program. Begun in 1998, the FFP1 program focuses on several “core capabilities” that are expected to improve controller decision-awareness and system performance in both the terminal areas and the en route airspace. This paper will focus on one such capability—a decision-support tool for en-route controllers known as the User Request Evaluation Tool (URET). More specifically, our aim is to determine the impact of URET on flight times. To do this, we employ statistical methods that compare changes in flight times before and after URET implementation for two sets of flights: one including flights going through sectors in which URET was implemented, and the second consisting of other flights.

The remainder of this paper is organized as follows. In the next section, we present general background on the FFP1 program and the intended contribution of this research to that program. Also, we offer explanation of URET capabilities and the development and deployment of the tool. Section 3 will discuss previous work on the benefits and impacts of URET, and Sections 4, 5 and 6 will discuss and explain the approach and methodology taken in this study. Discussion of some caveats is offered in Section 7, while conclusions and final remarks are offered in Section 8.

2. URET Functionality and Deployment

In collaboration with the aviation community, the Federal Aviation Administration has developed a plan for modernization of the National Airspace System (NAS), which is called the NAS Architecture. NAS Architecture defines methods for modernizing NAS in the 21st Century through application of new equipment, software, services, facilities, procedures and resources. An important aspect of this plan is a concept of operation known as Free Flight [5].

“Free Flight is a concept of air traffic management that permits pilots, dispatchers and controllers to share information and work together to manage air traffic” [7] from the planning phase and surface operations through en-route paths to arrival, without compromising safety. With Free Flight, use of prescribed flight routes based on ground-based navigation systems will be de-emphasized in favor of user-preferred routes. To achieve this, new technologies and procedures must be introduced.

A Free Flight Steering Committee was formed in 1995 to plan and oversee the implementation of recommendations for Free Flight. The committee was formed by Radio Technical Commission Aeronautics (RTCA) Task Force 3 to establish the strategy to implement Free Flight. RTCA “is a non-profit corporation formed to advance the art and science of aviation and aviation electronic systems for the benefit of the public” [12]. The committee suggested an incremental approach to the program, so that benefits could be realized as early as possible, and experience gained in the initial phases could guide the choices made in later ones.

Consistent with this approach, in 1998 a plan for developing the first Free Flight capabilities--called Core Capability Limited Deployment (CCLD)—was proposed. The first step in the development and deployment of these capabilities is Free Flight Phase 1 (FFP1), which includes five core capabilities:

- Surface Movement Advisor
- Collaborative Decision-Making
- Traffic Management Advisor
- Passive Final Approach Spacing Tool
- User Request Evaluation Tool (URET)

The FFP1 program is scheduled to end on December 31, 2002. After that date, the evolution to Free Flight will continue with Free Flight Phase 2, which will build upon the experiences of Free Flight Phase 1 and “introduce new capabilities from year 2003 through 2005. Free Flight Phase 2 has been chartered to geographically expand upon the successes of FFP1 as well as to conduct research to alleviate congestion and provide greater access to the NAS” [11].

With the exception of Collaborative Decision-Making, which has been implemented NAS-wide, FFP1 core capabilities have been deployed at a limited number of sites. The systems are constantly undergoing evaluation in order to refine them for future

deployments, and to determine whether, and where, such deployments should occur. The research presented in this paper is part of the evaluation effort for the URET tool.

Currently URET is in regular, daily use at the Indianapolis and Memphis Air Route Traffic Control Centers (ARTCCs). By the end of FFP1 URET will be deployed at five additional ARTCCs. These centers neighbor the Indianapolis and Memphis ARTCCs, as shown in Figure 1.

A URET prototype was initially demonstrated at the Indianapolis ARTCC in 1996. Since that time, it has been improved and updated regularly. Two “generations” of URET, called Build 1 and Build 2, are being implemented in Free Flight Phase 1. Build 1 includes the basic functions of the URET prototype and has been in use since the end of 1999. Build 2 includes functions that are considered necessary for operational acceptability at all seven Air Route Traffic Centers that are part of the FFP1, and is scheduled for deployment in 2002. Table 1 compares the functionality for Build 1 and Build 2.

2.1 URET Functionality

URET is a prototype of an automated conflict probe. Based on flight plans and actual flight tracks, the URET system models aircraft trajectories and predicts possible conflicts, up to 40 minutes into the future. It also allows controllers to check clearances for conflicts prior to issuing them to pilots, and to enter them directly into the Host Computer System (HCS). At the time of this writing, URET is deployed in the Indianapolis and Memphis Air Route Traffic Control Centers.

2.2.1 URET Build 1

URET Build 1 has been used in the Indianapolis Air Route Traffic Control Center (ZID) since January 1996 and in the Memphis Air Route Traffic Control Center (ZME) since Jun 1997. Below we overview the capabilities of Build 1 listed in Table 1 [5].

Trajectory Modeling

A core function of URET is to predict aircraft trajectories. To do this, URET processes real-time flight plan and track data from the main computer system, called the Host Computer System (HCS). The flight plan and the track data are combined with site adaptation data (active preferred routes, active altitude and speed restrictions), aircraft performance characteristics, and winds and temperature data obtained from the National Weather Service. These data are used to build four-dimensional flight trajectories for all flights in-bound to or within an ARTCC.

The trajectory modeling function includes three main tasks: route conversion, trajectory generation, and conformance bound determination. Route conversion is used for conversion of the flight plan route string into a series of coordinates. The coordinates are adjusted to conform to site-specific adaptation data. If the aircraft is already within the

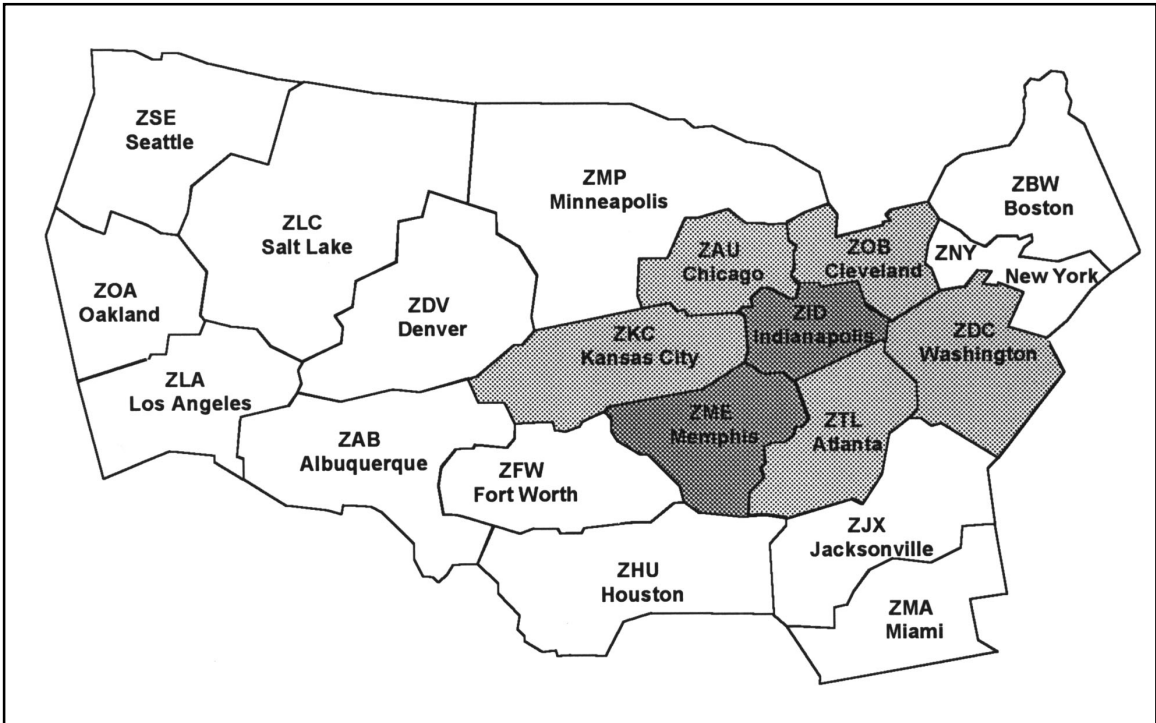


Figure 1. URET Sites During FFP1 (Source: FFP1 Performance Metrics Team Report, December 2000)

Table 1. URET Functional Capabilities

Functional Capabilities	Build 1	Build 2
Trajectory Modeling	X	X
Automated Problem Detection	X	X
Trial Planning	X	X
Two-Way Host Interface	X	X
Inter-facility	X	X
“Red Route” Processing	X	X
Arrival Stream Filters	X	X
Automatic Re-sectorization		X
Military Operations		X
Automated Coordination		X
Automated Re-plan		X
Hold Processing		X
ATC Preferred Route Processing		X

(Source: Free Flight Phase 1 Conflict Probe Operational Description)

facility and the current position data are available, then the route is further adjusted to take that into account.

Trajectory generation forecasts the predicted aircraft trajectory based on the already determined route, altitude transitions, and estimated flying times along the route. Times are estimated according to the planned aircraft speed, aircraft performance characteristics, forecast weather data (wind, temperature, and air pressure), and possible delays associated with the specific flight. The forecast includes the time and position of sector and center boundary crossings. This information is used for assigning the aircraft and its current flight plan to the appropriate sector, and is transferred to other centers that will be handling the flight.

Conformance bounds are determined to assess whether the modeled trajectory remains consistent with observed track. The conformance bound sets the maximum acceptable deviation of an aircraft position from its predicted trajectory. If the deviation exceeds this bound, which takes into account the navigational equipment of the aircraft, a new trajectory is built.

Track Management

The track management function of URET consists of conformance monitoring and re-conformance. Conformance monitoring constantly compares predicted trajectory position of aircraft to the reported position. If these positions are outside the previously calculated conformance bounds, the re-conformance function is invoked. The re-conformance function remodels the trajectory so that it again conforms to the reported aircraft location, speed, and course. Every time the re-conformance function updates the trajectory, the automated problem detection procedure, discussed below, is invoked to check for potential conflicts of the new trajectory.

Automated Problem Detection (APD)

This URET function detects and notifies the controller of a potential conflict. Both aircraft-to-aircraft and aircraft-to-airspace conflicts are considered. The aircraft-to-aircraft conflicts are probed up to twenty minutes into the future, while the aircraft-to-airspace (restricted airspace) conflicts can be probed up to forty minutes into the future. URET checks for conflicts every time a new flight plan is activated, and when any of the already active flight plans are amended or re-conformed.

When a conflict is detected, URET determines which sector to alert, and displays an alert to that sector (see the conformance boundaries function). Alerts are color coded in accordance with the severity of the conflict. Airspace alerts are coded blue; aircraft alerts are coded yellow or red. A yellow alert appears if the separation along with the conformance bounds between two aircraft is predicted to be below ATC minima. A red alert indicates a higher severity of the conflict—when horizontal separation between the aircraft themselves is predicted to fall below the minima. The amount of time between notification and predicted occurrence of a conflict is based on the type of conflict (red, yellow or blue) and the likelihood of conflict.

Trial Planning

Trial planning allows a controller to check a desired flight plan amendment for potential conflicts before issuing a clearance. If the desired amendment (trial plan) is conflict-free, the controller can send it to the Host Computer as a flight plan amendment. Amendments are initiated by the controller, either to resolve a potential conflict or to coordinate with a proposed amendment from another controllers, or by the pilot (requesting a change of route or altitude). Trial planning involves the trajectory modeling and conflict detection functions described previously.

Computer Human Interface

The computer human interface (CHI) includes text and graphic information. It consists of three displays. The aircraft list and plan display is text-based and manages the presentation of active flight plans, trial plans, and the conflict probe results. The graphic plan display provides a graphic view of aircraft routes and altitudes, predicted conflicts, and the trial plan results. The wind grid display shows the wind data overlaid on a sector map. The map includes sector and center boundaries, fixes, and wind data represented by arrows indicating wind direction, and a number indicating wind speed. Weather data are obtained from the National Weather Service, and includes temperature and air pressure as well as the wind data. The point-and-click interface of CHI enables quick entry and evaluation of the trial plan route, altitude, or speed changes, and sends these amendments to the Host Computer System.

2.2.2 URET Build 2

URET Build 2 adds capabilities to the URET Build 1 conflict probe. These should increase URET functionality and enable “seamless” operations and easier traffic transfer between different ARTCCs that will be part of Free Flight Phase 1 program.

Automated Re-plan

The Automated Re-plan function will allow a controller who has active control over an aircraft to create and evaluate trial plans for that aircraft periodically. That means the controller can create a trial plan for the aircraft, and if the trial plan is not problem-free, automated re-plan will continue to periodically check the status of the trial plan. When the plan is free of problems, the controller will be notified. This capability will relieve the controller from keeping track of the plan and its manual reassessment.

Automated Coordination

Automated Coordination will enable non-voice coordination between sectors that will be in consecutive control of a particular aircraft. The automated coordination makes it possible for a controller to send a trial plan to another sector in order to coordinate proposed amendments to the current flight plan. The trial plan sent for coordination is called the coordination plan. This plan can be deleted, timed-out, or accepted by the controller receiving it. It will be timed-out if there is no response to it after a certain amount of time has elapsed. If the coordination plan is accepted, the amendment to the current flight plan is sent to the Host Computer System, without need for controllers to verbally perform the coordination.

Hold Processing

The Hold Processing capability allows for placing an aircraft in holding, and will be applied manually when an aircraft reaches a planned holding area. The trajectory of the aircraft will not be probed for conflicts while it is in the planned holding area. However, all other aircraft will be probed for conflicts with the planned holding area. The planned holding area is represented as the rectangular airspace with upper and lower altitude boundaries. Conflicts are reported as aircraft-to-aircraft conflicts.

ATC Preferred Route Processing

ATC Preferred Route Processing checks for and notifies a controller when a flight is eligible for an ATC Preferred Route, based not only on the destination airport, but also on the adapted ATC Preferred Routes and eligibility criteria. This capability will also make it possible to create a trial plan with the insertion of the ATC Preferred Route.

Other Potential Enhancements

Other improvements to URET software are undergoing evaluation. Problem Analysis, Resolution and Ranking (PARR)[9] will provide a ranked set of problem resolution advisories for conflict and metering problems. These advisories will be offered to controllers in the form of URET trial plans to support strategic problem resolutions in a complex traffic environment. Additional improvements include the generation of conflict-probed sets of altitude, and possible direct-to-downstream fix maneuvers, as well as speed maneuvers. This will allow the controller to assess the conflict status of each of these maneuvers without having to manually check each element of the set for conflicts. Each altitude (or speed range) that is conflict-free appears in one color, while those that are not will appear in a different color. It is expected that these improvements will enhance safety, conflict resolution efficiency, user benefits, and conflict probe accuracy, as well as reduce the controller workload.

2.3 URET Development

As of February 2000, URET has been in daily use in the Indianapolis and Memphis ARTCCs. It took four years to reach this level of daily operational use, since the deployment occurred incrementally. This section describes this four-year process, which is also summarized in Figure 2.

URET development began in January 1995, when the FAA assigned the MITRE Corporation to develop a prototype of a conflict probe. The MITRE Center for Advanced Aviation System Development (CAASD) was given the task of conflict probe development, and the prototype was to be installed for evaluation in the Indianapolis ARTCC. CAASD developed the URET prototype based on the earlier conflict probe prototype – Automated En-Route Traffic Control (AERA). AERA was used for laboratory evaluations, with the participation of field controllers.

The first URET prototype was installed in Indianapolis ARTCC (ZID) in January 1996. URET software at the time included a “passive” conflict probe function, which allowed




Figure 2. URET Deployment.

conflicts to be displayed only upon a controller's request. At first the tool was used for operational evaluation. At that time, a mobile URET display was placed near the sector position for which it was being used. Evaluations began in February 1996, during which time three-person controller teams were working at the sector. The team consisted of an R-side (radar) controller, a D-side (strategic planner) controller, and a URET operator. In the initial evaluations, the URET operator was not supposed to communicate the conflicts detected by URET to the other members of the team. Since April 1996, however, such communication has been permitted [3].

By November 1996, a new version of URET that included the sector notification logic went online. In this version, only the sector in which the conflict is predicted to occur is notified. The probing for each sector extends to 20 minutes into the future. Also, the automated problem detection function was improved so that conflict notification time was determined according to the computed conflict probability. With this improvement, low probability conflict notification is delayed because in many cases such problems will resolve themselves, often as a result of trajectory re-conformance, before notification is necessary [3].

URET was installed at the Memphis ARTCC in June of 1997. It has been used on a daily basis at the Indianapolis ARTCC since October 1997, and at the Memphis ARTCC since November 1997. URET operations were scheduled for several hours every day. Beginning in 1998 an inter-facility capability, which enables data communication between neighboring centers, was added to the URET prototype. Because automated problem detection looks 20 minutes ahead, URET probes airspace that is outside the actual center border so that all possible conflicts can be detected. The outer border is drawn 200 nautical miles (nm) from the actual center boundary, and is known as the automated problem detection (APD) boundary. Each center has its own APD boundary. The inter-facility capability makes it possible for two centers to easily exchange data, especially data concerning conflicts likely to occur in the neighboring center. This ensures earlier controller notification on inbound traffic, and also increases controller awareness of possible problems that incoming traffic can create in the sector. The Indianapolis and Memphis URET systems are provided with direct digital two-way communications through the FAA's National Airspace Data Exchange Network (NADIN II). URET systems are using it for exchange of flight data between the two Host Computer Systems [2].

Another improvement—two-way communication between a controller and the Host Computer—was introduced in July 1999. This allowed controllers working on the URET workstation to send flight amendments directly to the Host Computer System with a click of the mouse. If the trial plan (created by a conflict notification or by pilot or controller request) proves to be problem free, the controller can send it to the Host Computer System as a flight plan amendment. The Host Computer automatically updates the flight plan and re-confirms the URET predicted trajectory of the flight. Any amendment “might change the sectors that the aircraft passes through, and amendment flight strips will be posted by the HCS to those sectors affected. In addition, the Aircraft List at affected

sectors will update to reflect the new amended flight plan, so that no communication is required with any intermediate sectors” [5]. This capability makes a controller’s job easier because the flight plans are updated automatically, which reduces the amount of coordination between sectors and the paper strip manipulations. Without the two-way communication capability, flight plan amendment required printing and distributing a new set of paper strips whenever the change would influence sectors beyond that where the amendment was issued. Two-way communication greatly expedites this process from the standpoint of the controller making the amendment, while at the same time expediting notification of controllers downstream that the change has taken place.

Beginning February 2000, the URET systems in both Indianapolis and Memphis began operating on a daily basis. URET is operational on all sector positions 22 hours a day, 7 days a week. It is still not possible to run 24-hour operations because the URET prototype needs a couple of hours to reinitialize, run analysis tools and record the previous day’s data for processing.

As noted above, additional improvements to URET are pending with the deployment of Build 2. Also, as part of the Free Flight Phase 1 program, URET should be operational in five more ARTCCs by the end of 2002.

3. Assessing URET—A Literature Review

Throughout its development, testing, and deployment, researchers have studied URET in order to determine whether it works in an operational sense, and to assess its realized and potential benefits. Results of these investigations are documented in many papers and reports. The URET research literature can be divided into two categories. One focuses on operational evaluation of the tool itself, while the other considers benefits that the tool could provide or is already providing. This section summarizes previous studies on URET, with particular attention to work in the latter category.

Earlier papers on URET were primarily concerned with the field evaluation of the tool, based on feedback from controllers, and assessment of its functional performance. A main concern was how controllers would accept URET and its capabilities. In Brudnicki [4], results from the initial field evaluations of URET were discussed, as well as the controllers' opinions on its reliability. After the initial training and evaluation, participating controllers had a very clear consensus "that the prototype capabilities are sufficiently accurate and suitable to support conflict detection, conflict resolution and planning at the D-position" of the sector team. They also stated that, "URET capabilities should be implemented at the sector as a D-position tool and should replace the paper flight strips." Further evaluations were focused on conflict probe functional performance and its quantification. Some inaccuracies in software were uncovered and corrected [2]. "Overall, the approach of combining the lab analysis of real-world data with field trial results provided a useful mechanism to validate the core URET functions and to refine them to the level of being operationally acceptable for problem solving and planning at the sector" [3].

Papers by Arthur [1] and Kirk [9] focus on the process of defining performance metrics to be used to compare different URET versions and URET prototype enhancements. The incremental approach produced different versions of URET, and it was necessary to find a way to compare their functionality. Each system upgrade was evaluated for its technical performance and for areas that needed improvement.

Research on URET benefits focused on three areas [5]: safety, controller productivity, and benefits to airspace users. URET capabilities that can increase safety include Automated Problem Detection and Trial Planning. These help in preventing controller operational errors. "An operational error is an event where two aircraft under positive control come closer together than the required separation standard (5 nm laterally and 1000 or 2000 ft vertically for en route radar control) because of an error on the part of the ATC system" [2]. Analysis of the probable effect of URET on operational errors was discussed in Kerns [2] and Celio [5]. Applying the tool retrospectively in situations where operational errors occurred, it was shown [2] that for operational errors where two aircraft were on conflicting paths for some time before the conflict occurred, "URET generated conflict alerts with more than adequate time for controllers to take action to avoid the operational error." For operational errors that were the result of ill-advised clearance, URET couldn't provide a timely alert, but with the Trial Planning function

controllers could check the clearance prior to issuing it. It was also found that URET performs well for a wide range of traffic densities, including very high and low densities, which pose the greatest challenges for human controllers. The ability of the tool to prevent operational errors depends on timely updating of vectoring clearances into the Host computer.

It is often stated that URET enables controllers to handle more aircraft without a proportional increase in controller workload, thus improving controller productivity. The National Research Council panel on human factors in air traffic control automation stated the following concerning flight data management workload: "Aggregate observations suggest that an interactive integrated display or interface that provides more direct access to both flight and radar data could enhance controllers' performance without a reduction in situation awareness" [5]. However, there is little evidence that URET enables controllers to handle increased traffic. Part of the problem is that URET is intended for a D-side controller while the R-side controller still does not have such a tool available.

In addition to its safety and productivity benefits, URET is expected to improve the quality of ATC services provided to users. Specifically URET may allow greater use of preferred flight trajectories by affording greater flexibility for users to choose routes, altitudes, and speeds. The National Airspace System today is highly restrictive in these respects. The restrictions help controllers manage traffic without compromising safety. There are several types of restrictions [7]:

- Preferred IFR High and Low Altitude Routes that must be flown between certain airports or through certain air spaces;
- Standard Terminal Arrival Routes and Standard Instrument Departures that specify the route to be flown for arriving at or departing from airports;
- Altitude restrictions for aircraft arriving at or departing from specific airports or airspace;
- One-way airways in congested traffic areas;
- Altitude-for-direction rules in determining allowed cruising altitudes;
- Dynamic restrictions applied by traffic management when conditions exist preventing a normal flow of traffic in an airspace, such as miles-in-trail, ground delays, and ground holds;
- Lower cruise altitudes than those requested.

Several papers [5,8,4] convey research results concerning relaxation of some or all of the imposed restrictions. One study examined potential flight-time savings and associated monetary benefits of the ultimate free-flight environment [5], using several simulation scenarios. Another assessed the excess mileage flown on airways compared to direct routings between origin and destination (the ultimate goal of free flight). Research was also conducted on savings that can be achieved by removing altitude restrictions, which would enable more fuel-efficient flight profiles [8]. All of these studies showed that benefits obtained amount to several hundred million dollars per year [5].

The focus of these earlier studies, which were conducted in parallel with the URET prototype development, was to identify benefit mechanisms and predict their magnitude.

Later work, in contrast, has investigated the actual consequences of URET implementation. Since February 1999, the FAA's Free Flight Phase 1 office has been monitoring URET benefits and utilization. This work concentrates on changes observed after July 1999, when two-way communication between URET and the Host Computer System was implemented. As explained above, this is considered a critical date because with two-way communication, the workload from modifying flight plans was greatly reduced.

The Free Flight Phase 1 Metrics Team publishes semiannual reports on all Free Flight Phase 1 automation tools [10,11]. There are several metrics concerning URET and the change in distance caused by URET: distance saved for lateral amendments, excess distance, and en route distance.

The first of these metrics is based on all lateral amendments made in ZID and ZME. A lateral amendment is a change of the aircraft route without altitude change. "The metric measures the average daily sum of nautical miles (nm) changed as the result of an amendment; i.e., the distance from the point of the amendment to the destination airport. It includes all lateral amendments entered into the Host Computer System for the specified time, not only URET amendments" [10]. The purpose of this metric is to show whether URET is enabling aircraft to fly shorter routes. For calculation of this metric only the busiest hours of the two busiest days of the week are used. The metric shows an increase from approximately 500 nm average daily savings (per center) to more than 4000 nm average daily savings in distance [10] from May 1999 to May 2000. These measurements are based on simple geometric distances, without adjustment for wind effects.

The excess distance metric is calculated as the difference between the actual distance flown across the center and the great circle distance between center entry and exit points. This metric is calculated for ZID, ZME, and for the additional five URET sites that are to be part of the Free Flight Phase 1 program. The purpose of the metric is to determine the impact of URET on the length of the path flown within the center. The metric shows an increase between May 1999 and May 2000 instead of the expected decrease. For ZME it increased from "slightly less than 4 nm per aircraft to slightly over 4 nm per aircraft; at ZID the increase was from about 5.5 to 5.8 nm per flight".

Both of the above metrics focus on the benefits for flights that are traversing URET prototype sites and that are accrued within the centers. A third, the en-route distance metric, "explores this distance savings question by looking at the entire 'en-route' portion of a flight." The metric is calculated as the average en-route distance for each of the ten chosen city-pairs, for "each of the selected analysis days" [11] (no explanation is offered on the selection criteria). The analysis shows a slight decline in distance between all selected city-pairs. The results are statistically significant for city-pairs that traverse ZME, while they are not statistically significant for the ZID city-pairs, as shown in Figures 3 and 4.




Figure 3. En-Route Distance Trend, ZME (Source: FFP1 Metrics Report, June 2001).

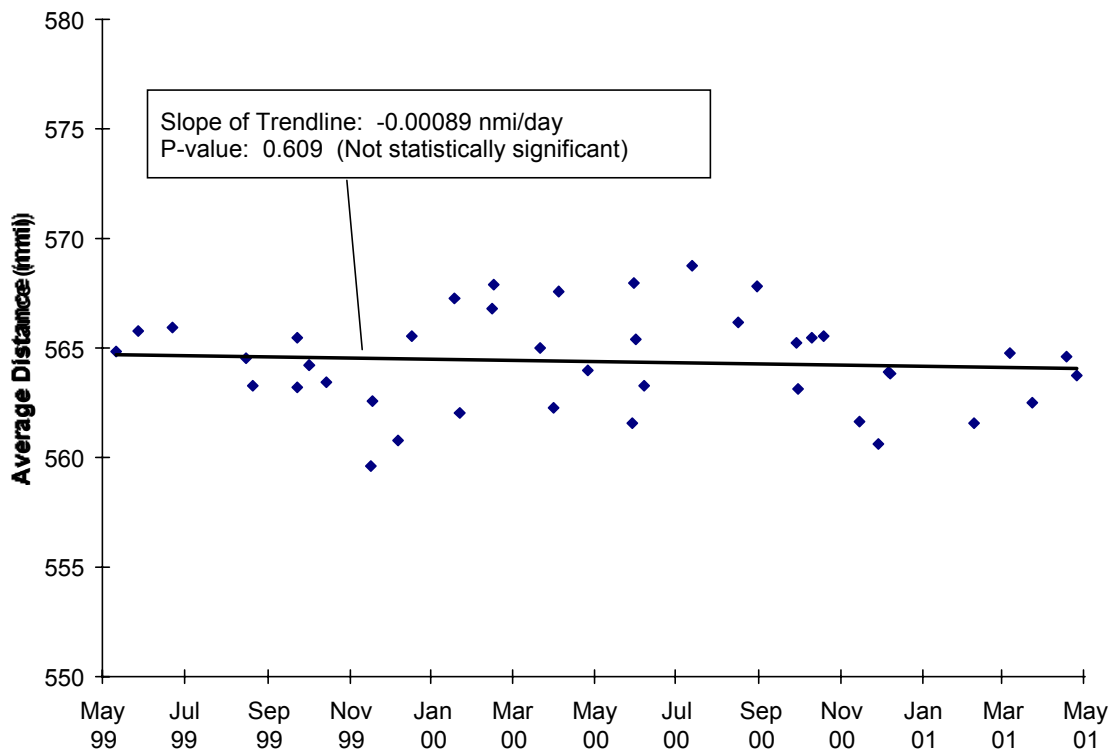


Figure 4. En-Route Distance Trend, ZID (Source: FFP1 Metrics Report, June 2001).

In conclusion, there has been considerable research on both the functionality and benefits of URET. Even so, there are major gaps in our knowledge, particularly on the benefits of URET to airspace users. Most of the benefits studies are either concerned with projecting the ultimate impact of the free flight environment, or, when considering realized benefit from URET implementation, with estimating the local impact on distance flown within URET airspace. It is not possible to conclusively determine the impact of the tool by considering only the airspace in which the tool is used. One way of improving the assessment is to broaden the scope of analysis, by comparing outcomes for entire flights that do and do not use URET airspace. And in so doing, it is better to focus on metrics based on flight time than on distance flown. This permits assessment of benefits that accrue both within and outside the URET airspace, to take into account the effects of winds, and to consider effects related to congestion and throughput as well as the directness of routing. The remainder of this report documents such an approach.

4. Methodology Overview

We now turn to our own study of the effects of URET. There are three distinguishing features of our approach. First, we adopt a quasi-experimental method in which we compare flights that use URET sectors to those that do not. Thus, instead of before-after comparisons or trend analyses for flights that use the URET centers, we compare changes for such URET flights to changes for flights that do not use these centers. The latter are used as a control group, and the former as the treatment group in our quasi-experiment. This approach eliminates effects of NAS-wide performance trends that could be confused with the impacts of URET in a simple before-after analysis. For example, there have been significant changes in air traffic management procedures over the past several years as a result of collaborative decision-making (CDM).

Second, we focus on flight time metrics rather than distance metrics. It is flight time, rather than distance per se, that is of economic significance when considering the value of URET to users of NAS. While changes in distance flown are likely to correlate strongly with flight time changes, the effects of winds and congestion makes this correlation less than perfect.

Third, we focus exclusively on “end-to-end” metrics rather than metrics based on portions of flights within the ZID and ZME centers. Again, these are the metrics of ultimate economic significance: there is no gain from a shorter flight time on a particular portion of a flight unless it effects the overall flight time. Moreover, there are a variety of ways in which the presence of URET could influence flight times outside the ZID and ZME centers. For example, if the capacities of sectors in these centers were increased, then there could be a reduction in miles-in-trail restrictions and thus of flight times in the “upstream” centers. In addition, there could be flights that avoided these centers altogether in the pre-URET period by taking more circuitous routes.

We take two different approaches, both with the above features, to assess the impact of URET. In the first approach the unit of analysis is a city-pair. We compare changes in average flight times between two sets of city-pairs—one set that is generally routed through the URET centers, and one set that is not. In the second approach the unit of analysis is the individual flight. Using sets of flights that do and do not use the URET centers, and which take place before and after URET implementation, we estimate the effect of URET implementation on the URET flights. In the following sections, we describe the methodologies and present the results of these analyses in turn.

5. Average Flight Time Analysis

5.1 Methodology

In this analysis, the metrics chosen are based on average flight times for two sets of city-pairs. We define URET city-pairs as those that generally use URET centers (ZID and ZME) in their routings, and non-URET city-pairs as those that do not use these centers. (More precise definitions are provided later on.) Our task is to determine how average flight times for URET city-pairs changed relative to those for non-URET pairs, comparing periods before and after a significant milestone in URET implementation. More formally, we test the null hypothesis:

H₀: Flights times for URET and non-URET city-pairs changed the same amount between pre- and post-milestone periods.

Data used to obtain this analysis came from two sources: the Airline Service Quality Performance (ASQP) database and excerpts from Enhanced Traffic Management System (ETMS) data for ZID and ZME. The ASQP database contains information on all flights performed by the ten biggest passenger airlines in the US. The database contains data on the scheduled departure and arrival times, actual departure and arrival times, taxi-out and taxi-in times, wheels-off and wheels-on times, and various time intervals between these times (such as arrival delay against schedule). ASQP has no data on route flown, however. Thus, in order to identify the flights that were traversing URET airspace, boundary-crossing data derived from ETMS is used. The boundary crossing data contains information for each flight that crosses ZID and ZME center boundaries, such as flight time and flight distance within these centers.

The time periods used for the analysis were February to July, 1999 and 2000. Data for 1999 is used as a benchmark, because URET was still not fully operational in that period. As explained above, the most important milestone in URET deployment was the initiation of two-way communication between URET “stations” and the Host Computer in early July of 1999. After that time the utilization of URET increased dramatically. Also, in February 2000, daily use of URET in ZID and ZME began. Therefore, the months from February to July 1999 are defined as the “before” period, while the same months in 2000 are identified as the “after” period. Before-after comparisons are made for individual months (February 2000 and February 1999, et cetera) and for the entire six-month period.

To identify URET and non-URET city-pairs, we merged the ASQP data with the ETMS boundary crossing data to determine, for each city-pair, the proportion of flights using ZID and ZME. (We considered a flight to use these centers if they spent more than 10 minutes in either one or both. We define our city-pairs directionally so that the A-B pair and the B-A pair are considered to be distinct.) We classified each city-pair from flights in the ASQP database as follows (see also Figure 5):

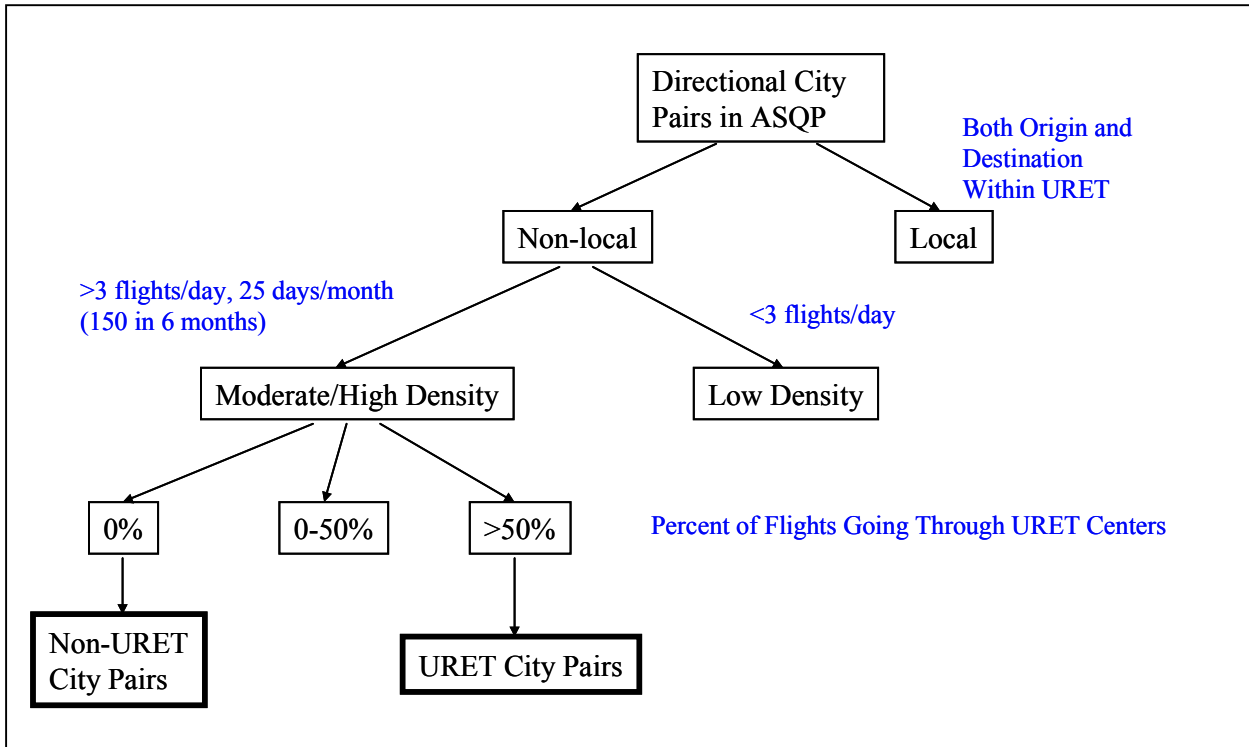


Figure 5. URET and Non-URET City-pairs Generation.

1. Low traffic density city-pairs, where the number of flights is less than three per day (for more than 5 days a month, or 30 days for the 6 months period);
2. Local city-pairs, where both the origin and destination airports are in the ZID or ZME regions, and that are not in Class 1;
3. Non-URET city-pairs, flights between which do not use URET airspace at all in either the before or after period and that are not in Class 1;
4. URET city-pairs for which 50% or more of the flights go through URET airspace in both the before and after periods, and that are not in Class 1;
5. All others, which include those with more than 0% but less than 50% of flights using URET airspace in either the before or the after period, and that are not in Class 1.

Class 3 and 4 city-pairs are used in the analysis while the others are not. The first class is not considered because of the sample size considerations. The second is eliminated because those flights generally do not use URET airspace during the en-route portions of their flights. The fifth class is excluded because, if a large fraction of flights for a city-pair go around the URET airspace, we have less confidence that a change in the average flight time for the city-pair is due to URET.

Two metrics are calculated for the URET and Non-URET city-pairs. The Average Airborne Time Change (AATC) is based on the time between wheels-off and wheels-on, which we term the airborne time (AT). The Average Flight Time Change (AFTC) is based upon the flight time (FT), which we define as the total of the departure delay, taxi-out time, and airborne time. Unlike AATC, the AFTC includes the time a flight spends at its origin airport. URET could affect this if there is ground holding resulting from en-route traffic congestion.

Figure 6 shows a flowchart that represents calculation of these metrics. The AATC and AFTC are calculated in the same manner. In the following text we illustrate the procedure using the former. The AATC metric is calculated as:

$$AATC_i = \frac{\sum_{k \in i} AATD_k}{N_i} \quad (1)$$

where:

$AATC_i$ is the average airborne time change of city-pair class i (URET or non-URET);
 $AATD_k$ is the difference in average airborne times for 1999 and 2000, for city-pair k ;
 N_i is the number of city-pairs in class i .

The difference in average airborne times for a given city-pair k , $AATD_k$, is calculated as

$$AATD_k = AAT_{99k} - AAT_{00k} \quad (2)$$

where:

AAT_{99k} is an average airborne time for city-pair k during February-July of 1999;
 AAT_{00k} is an average airborne time for city-pair k during February-July of 2000;




Figure 6. Metrics Calculation Flow-chart.

To calculate the average airborne time for a given time period t (either February-July 1999 or February-July of 2000) and city-pair k , AAT_{tk} , we first calculated a daily average and then averaged over the days in the period:

$$AAT_{tk} = \frac{\sum_{d \in D_{tk}} \frac{\sum_{f \in F_{dk}} AT_f}{M_{dk}}}{N_{tk}} \quad (3)$$

where:

- AAT_{tk} is the average airborne time for city-pair k and time period t ;
- AT_f is the airborne time for flight f ;
- D_{tk} is the set of days included in the average for city-pair k and time period t ;
- N_{tk} is the number of days in the set D_{tk} ;
- F_{dk} is the set of flights included in the daily average for day d and city-pair k ;
- M_{dk} is the number of flights in the set F_{dk} .

The set of flights included in the daily average (i.e. F_{dk}) includes any completed flight with an arrival delay of less than three hours. In determining which days to include in the averaging (i.e. the days included in D_{tk}) we excluded those days where the set F_{dk} contains less than three flights. The number of days excluded cannot be more than 5 for the month-by-month analysis, or more than 30 days for the 6-months analysis. If that number of excluded days exceeds 5 (30), then that city-pair belongs to the group 1 as defined above and is not considered in the analysis. Table 2 shows the airports used as origins and destinations for URET city-pairs, as well as their frequencies of occurrence as an origin (or destination) in different URET city-pairs. Table 3 shows the same for non-URET city-pairs

5.2 Results

To compare the airborne and flight time changes for URET and non-URET city-pairs we use t-tests. The t-test is applied to the AATC and AFTC metrics to determine if the means for the two populations (URET and non-URET city-pairs) are statistically different. In other words, we test the null hypotheses:

$H0(1)$: $AATC_{URET}$ and $AATC_{non-URET}$ are equal.

$H0(2)$: $AFTC_{URET}$ and $AFTC_{non-URET}$ are equal.

Table 4 presents the results for the AATC analysis performed on a month-to-month basis, as well as for the entire six-month period (in bold). Table 5 does the same for the AFTC analysis. The last column of each table presents the significance level of the hypothesis test. If the p-value is low, then it is unlikely that the observed difference in flight time change would occur if the null hypothesis were true. In other words, we need to reject the null hypothesis. Figures 7 and 8 depict graphically the results in Tables 4 and 5 respectively.

Table 2. Airports in URET City-pairs and Number of Occurrences

Airport	Number of Occurrences		Airport	Number of Occurrences	
	As Origin	As Destination		As Origin	As Destination
ABO	1	1	I.GA	8	7
ALB	1	2	LIT	4	4
ATL	40	35	MCI	9	10
AUS	2	4	MCO	13	13
BDL	3	3	MDW	4	5
BHM	7	4	MEM	24	24
BNA	17	21	MIA	7	7
BOS	4	4	MKE	3	3
BTR	2	2	MSP	9	9
BUF	2	1	MSY	7	10
BWI	13	10	OKC	2	2
CLE	8	8	OMA	1	1
CLT	14	16	ORD	22	17
CMH	14	14	ORF	3	3
CVG	37	38	PBI	1	1
DAL	1	1	PDX	2	1
DAY	3	4	PHL	14	13
DCA	7	5	PHX	6	5
DEN	9	9	PIT	16	14
DFW	29	28	PWM	1	1
DTW	18	18	RDU	5	1
EWR	9	6	RIC	2	5
FLL	4	3	ROC	1	3
GPT	1	1	RSW	1	1
GRR	1	1	SAN	2	2
GSO	1	2	SAT	1	2
GSP	1	1	SBN	1	4
HOU	4	6	SDF	8	1
HSV	1	1	SEA	2	8
IAD	7	8	SFO	5	1
IAH	12	15	SLC	3	6
IND	10	12	STL	37	2
JAN	4	4	TPA	8	3
JAX	3	3	TUL	1	30
JFK	3	2	VPS	1	8
LAS	2	3			1
LAX	9	11			2
LEX	1	1			1

Table 3. Airports in Non-URET City-pairs and Number of Occurrences

Airport	Number of Occurrences		Airport	Number of Occurrences		Airport	Number of Occurrences		Airport	Number of Occurrences	
	As Org	As Des		As Org	As Des		As Org	As Des		As Org	As Des
ARF		7	FAI	1	1	MHT	3	3	SFO	23	20
ABQ	10	10	FAR	1	1	MIA	8	11	SGF	1	0
AGS	1	1	FAY	1	1	MKE	3	3	SHV	1	0
ALB	4	3	FLL	11	12	MLB	1	1	SJC	14	13
AMA	2	2	FNT	1	1	MLI	1	1	SJU	8	8
ANC	3	3	FSD	1	2	MOB	1	1	SLC	25	23
ATL	22	25	GEG	5	5	MOT	1	1	SMF	12	12
AUS	6	5	GRB	1	1	MSN	2	1	SNA	12	13
AVL	1	1	GRR	3	1	MSP	28	1	SRQ	2	1
AVP	1	0	GSO	2	3	MSY	5	26	STL	13	14
BDL	5	7	GSP	2	3	MYR	1	4	SWF	0	1
BHM	1	1	HNL	2	2	OAK	12	13	SYR	4	4
BIL	2	2	HOU	7	2	OGG	1	1	TLH	1	1
BIS	1	1	HRL	1	8	OKC	6	6	TPA	6	10
BOI	7	7	HSV	1	1	OMA	6	5	TRI	1	1
BOS	14	12	IAD	7	5	ONT	10	10	TUL	3	3
BTR	1	0	IAH	16	19	ORD	21	20	TUS	5	4
BTV	1	1	ICT	2	1	ORF	3	2	TYS	2	2
BUF	6	5	ILM	1	1	PBI	4	4			
BUR	6	6	JAX	3	5	PDX	17	17			
BWI	9	9	JFK	6	4	PHL	25	25			
CAE	2	2	JNU	1	1	PHX	29	29			
CHS	2	2	KTN	1	1	PIT	13	16			
CID	2	2	LAN	1	1	PNS	2	2			
CLE	8	6	LAS	23	22	PSC	1	1			
CLT	29	30	LAX	20	19	PVD	5	4			
COS	3	2	LBB	2	1	PWM	1	1			
CRP	1	1	LGA	11	10	RDU	6	6			
DAB	1	1	LGB	2	2	RIC	4	3			
DAL	10	10	LNK	1	0	RNO	11	11			
DCA	8	8	MAF	1	1	ROA	1	1			
DEN	27	27	MBS	2	2	ROC	3	3			
DFW	24	28	MCI	10	9	RST	1	1			
DSM	3	3	MCO	14	15	RSW	1	2			
DTW	20	23	MDT	3	1	SAN	13	14			
ELP	11	10	MDW	1	5	SAT	8	9			
ERI	1	1	MFE	2	2	SAV	2	1			
EUG	1	1	MFR	1	1	SBA	1	1			
EWR	15	11	MGM	0	1						

Table 4. AATC Metric Values for URET and Non-URET City-pairs

Time Period	AATC _{URET} (min)	Number of City-pairs	AATC _{non-URET} (min)	Number of City-pairs	Difference	P-Value
Whole period	-0.248	539	-0.253	859	0.005	0.9530
February	-0.243	457	-0.322	962	0.079	0.4627
March	-0.062	518	-0.203	1071	0.141	0.2530
April	-0.005	516	-0.180	997	0.174	0.2585
May	-0.325	490	-0.204	985	-0.121	0.4509
June	-0.145	498	-0.196	1013	0.051	0.7081
July	-0.589	497	-0.394	1040	-0.196	0.1212

Average Airborne Time Improvement




Figure 7. Average Airborne Improvement.

Table 5. AFTC Metric Values for URET and Non-URET City-pairs

Time Period	AFTC _{URET} (min)	Number of City-pairs	AFTC _{non-URET} (min)	Number of City-pairs	Difference	P-Value
Whole period	-1.095	539	-1.679	859	0.585	0.0042
February	-2.189	457	-2.837	962	0.648	0.0065
March	-0.530	518	-0.797	1071	0.268	0.3230
April	0.203	516	-0.745	997	0.947	0.0004
May	-1.227	490	-1.124	985	-0.103	0.7479
June	-1.969	498	-3.177	1013	1.208	0.0018
July	-0.433	497	-1.030	1040	0.598	0.1497

Average Flight Time Improvement




Figure 8. Average Flight Time Improvement.

Table 4 suggests that the null hypothesis $H0(1)$ cannot be rejected for the six-month period as a whole or for any month alone. All of the AATC values in the table are negative, which means that on average airborne times were lower in 1999 than they were in 2000 for both the URET and the non-URET city-pairs. Even though the AATC values for URET city-pairs were better than the ones for non-URET city-pairs for most of the months and the whole period, p-values suggest that these differences could be simply due to chance. In sum, this analysis yields little or no evidence that airborne times decreased as a result of URET.

Table 5 shows AFTC metric calculated across the whole period, as well as for each month. From these results we can again see that most of AFTCs (excepting those for April) are negative, suggesting an overall degradation in performance in 2000 compared to 1999. However, the AFTC values computed for the whole period, and most of the months, show that Average Flight Times for URET city-pairs worsened less than the ones for non-URET city-pairs. Moreover, p-value suggests that we should reject null hypothesis $H0(2)$ that the changes in average flight times are equal for both groups of city-pairs. In other words, the AFTC analysis gives strong evidence that Average Flight Times for URET flights decreased relative to those for non-URET flights in the year after URET implementation. The average decrease is approximately 0.6 minutes per flight.

In sum, the average flight time analysis provides three findings. First, it gives no evidence that average airborne times changed as a result of URET. Second, the analysis strongly suggests that average flight times decreased in conjunction with URET implementation. Third, from the previous two findings, we must conclude that the change in average flight time was in the ground time component rather than the airborne time component.

The analysis presented in this section is based on an aggregate approach. It requires averaging over individual flights, and designating them as URET and non-URET based on the city-pair rather than the actual flight itinerary. This simplifies the statistical analysis, but entails a substantial loss of information from the original data. In the next section, we analyze the data in a different way—one that entails a more complicated model, but which allows analysis of individual flights rather than city-pair averages.

6. Individual Flight Time Analysis

In the previous section, we investigated the flight-time impact of URET implementation by comparing average flight time changes for URET city-pairs and non-URET city-pairs. In this section, we report on a more disaggregate analysis, based on individual flight times instead of city-pair averages. The disaggregate approach vastly increases the amount of information going into the analysis, but also requires a more complex analysis procedure. On balance, however, the disaggregate approach appears to be more successful in capturing the benefits of URET.

6.1 Individual Flight Times Model

For this analysis, we estimate a model of flight times of individual ASQP flights. As before, we measure flight time as:

$$\text{Flight Time} = \text{Departure Delay} + \text{Taxi-out Time} + \text{Airborne Time} \quad (4)$$

Airborne time, as before, is the time aircraft spends between wheels-off the runway and wheels-on the runway. Departure delay is the time between the scheduled departure time and the time the aircraft leaves the gate. It may be positive or negative. The taxi-out time is the time between leaving the gate and wheels-off runway time. We estimate models for the total flight time and its individual components.

With the individual flight times model we must control for the influences of several factors on flight times, including:

- Distance between origin and destination,
- Direction of flight, which captures the winds aloft effect,
- Airport fixed effects,
- Overall trends in flights times, and
- Effect of URET implementation.

Distance is calculated as the great circle distance between origin and destination airports. The distance is modeled by a piece-wise linear function of distance versus time, as will be shown in Figure 9. Use of the piece-wise linear function allows differences in flight speed for different flight phases to be captured. The different ranges in the function are used as variables in the model. Ranges are: 0-200 nm, 200-500 nm, 500-1000 nm and over 1000 nm. The first range, 0-200 nm captures the influence of the initial climb, practically the part of flight that is under the terminal control jurisdiction. Climb to cruise altitude is represented by 200-500 nm range. To illustrate how these ranges are used, a flight of 900 nm would include 200 nm in the first range, 300 in the second, 400 in the third, and 0 in the fourth range.

To capture the effect of flight direction, we include in our model the differences in latitude and longitude between flight origin and flight destination. Thus, a flight from the northeast to the southwest would have a negative latitude change and a positive longitude change. We expect flight times to increase with the longitude variable because the

prevailing wind direction is westerly. The latitude variable, in contrast, is not expected to have a strong effect.

Airport congestion and delay, along with airfield differences that affect taxi times, are expected to be a significant source of variation in flight times. In this study, we control for these kinds of systematic inter-airport differences by using airport fixed effects. Separate fixed effects for arriving and departing flights are included. The effects are captured using airport dummy (0-1) variables. To illustrate, a flight from BOS to SFO would include a fixed effect for BOS departures and for SFO arrivals; the BOS departure dummy variable and the SFO arrival dummy variable would both be set to 1. All other airport dummy variables would be set to 0 for this flight. Our model includes fixed effects for the 40 US airports reported to have the highest delays in the 1998 FAA Airport Capacity Enhancement Plan [11].

The distance, directional, and airport fixed effects are included to control for major sources of flight time variation that are unrelated to URET implementation. As such, they are “nuisance” parameters that must be estimated in order to isolate the effect of primary interest. To estimate this effect, we include three additional dummy variables. First, there is a URET airspace (ZID and ZME) variable that is set to 1 if a flight spends more than 10 minutes in a URET sector. Second, we include a post-implementation dummy variable that is set to 1 for flights in the February-July, 2000, time period. Our data set is restricted to flights from this period and from the corresponding period one year earlier. Finally, we include an interaction dummy variable that indicates that a flight uses URET airspace in the post-implementation period.

It is the last of these variables that captures the effect of primary interest. The URET airspace variable captures persistent effects of flying through ZID and ZME—those that occur in both the pre- and post-implementation periods. These may result from the air route system, winds, weather, or congestion. The post-implementation variable controls for systematic differences between pre- and post-implementation flight times that exist for all flights, whether or not they fly through URET airspace. Possible sources of such differences include growing congestion, changes in air traffic management procedures, and meteorological factors. The interaction variable, in contrast, captures post-implementation changes in flight times that occur only for flights that go through URET airspace. The most likely source of these changes is the implementation of URET itself.

Thus, the specification for the individual flight time model form is as follows:

$$\begin{aligned}
 FlightTime = & \tau + \sum_{\ell} \alpha_{\ell} L_{\ell} + \beta_{lat} \cdot X_{lat} + \beta_{lon} \cdot X_{lon} + \sum_{i=1}^{40} [\delta_{ai} \cdot A_i + \delta_{di} \cdot D_i] + \\
 & \mu \cdot URET + \pi \cdot AFTER + \theta \cdot AFTER \cdot URET
 \end{aligned} \tag{5}$$

where:

L_{ℓ} is the distance flown in distance range ℓ (ranges are 0-200, 200-500, 500-1000, and over 1000 nm)

X_{lat}	is the destination latitude minus the origin latitude
X_{lon}	is the destination longitude minus the origin longitude
A_i	is a dummy variable set to 1 if airport i is the arrival airport, 0 otherwise
D_i	is a dummy variable set to 1 if airport i is the destination airport, 0 otherwise
$URET$	is a dummy variable set to 1 if the flight is in URET airspace (ZID or ZME, or both) for more than 10 minutes
$AFTER$	is dummy variable set to 1 if the flight took place in year 2000, the after period of the analysis (1999 being before)
$AFTER*URET$	is a dummy variable set to 1 if the flight took place over URET airspace in the after period of analysis
$\tau, \alpha_\ell, \beta_{lat}, \beta_{lon}, \delta_{ai}, \delta_{di}, \mu, \pi, \theta$	are coefficients to be estimated

The model was estimated on six different data sets. A given data set includes all flights in the ASQP database for corresponding months of the February-July 1999 and February-July 2000 time periods (February 1999 and February 2000, March 1999 and March 2000, et cetera) that met two additional criteria. First, the flight must have reached its scheduled destination within three hours of the scheduled arrival time. Second, the flight origin and destination must be in the continental U.S. The month-by-month partitioning allows coefficient values to vary seasonally, in response to monthly differences in weather conditions and demand.

6.2 Model Estimates Discussion

As explained above, our primary interest is in the estimates for θ ($AFTER*URET$ variable coefficient). Before turning to that, however, we present the entire estimation results for the June 1999-2000 data set, which is generally representative. Table 6 shows estimation results for all of the variables, except the airport fixed effects, for the month of June. Table 7 shows estimates of coefficients for airport fixed effects for the month of June.

Table 6 shows that all coefficient estimates are statistically significant. It can also be seen that the signs coefficients have the expected signs and reasonable magnitudes. The estimates for the distance variables are all positive, as common sense suggests they should be. Moreover, there is a clear negative correlation between the distance coefficient and its associated distance range. At the lowest range, which corresponds to flight operations near the terminal, the coefficient implies an average ground speed of 350 kts. At the 1000+ range, this increases to 460 kts. Figure 9 represents the piece-wise linear function used in the model, calibrated on the estimates shown in Table 6.

The effect of longitude difference is, as expected, positive, implying that westbound flights take longer than eastbound ones (Figure 10). The magnitude of the latitude effect is much smaller, also as expected, but statistically significant. Going north takes slightly longer than going south, a result for which there is no obvious explanation.

Table 6. Individual Flight Time Coefficients for June

Coefficient	Description	Estimate	Standard Error	P-Value
τ	Intercept	24.470	0.3672	<0.0001
α_1	Distance in 0-200 nm range	0.169	0.0020	<0.0001
α_2	Distance 200-500 nm range	0.146	0.0005	<0.0001
α_3	Distance 500-1000 nm range	0.139	0.0003	<0.0001
α_4	Distance 1000+ nm range	0.131	0.0002	<0.0001
β_{lat}	Difference in latitude	3.426	0.5969	<0.0001
β_{lon}	Difference in longitude	21.890	0.2274	<0.0001
μ	<i>URET</i> dummy	4.172	0.0854	<0.0001
π	<i>AFTER</i> dummy	-1.322	0.1257	<0.0001
θ	<i>AFTER URET</i> interaction	-3.260	0.1506	<0.0001
Adjusted R ²		0.8168		
Number of Observations		840656		

Table 7. Airport Fixed Effects Coefficients for June

Origin	Estimate	P-Value	Destination	Estimate	P-Value
JFK	22.390	<0.0001	JFK	11.076	<0.0001
EWR	22.348	<0.0001	EWR	13.657	<0.0001
LGA	21.915	<0.0001	LGA	9.328	<0.0001
PHL	19.928	<0.0001	PHL	10.474	<0.0001
ORD	18.742	<0.0001	ORD	9.540	<0.0001
STL	16.412	<0.0001	STL	6.066	<0.0001
DFW	15.129	<0.0001	DFW	2.284	<0.0001
BOS	13.759	<0.0001	BOS	6.768	<0.0001
IAD	13.119	<0.0001	IAD	5.817	<0.0001
DTW	12.109	<0.0001	DTW	-0.024	0.9106
ATL	11.154	<0.0001	ATL	4.579	<0.0001
MDW	10.720	<0.0001	MDW	1.541	<0.0001
IAH	9.883	<0.0001	IAH	2.881	<0.0001
MIA	9.362	<0.0001	MIA	5.546	<0.0001
MSP	8.948	<0.0001	MSP	0.328	0.1591
CTL	8.413	<0.0001	CTL	0.784	0.0009
PIT	7.678	<0.0001	PIT	0.556	0.0309
CVG	7.166	<0.0001	CVG	3.597	<0.0001
BWI	6.402	<0.0001	BWI	5.193	<0.0001
LAS	6.043	<0.0001	LAS	1.535	<0.0001
SFO	5.966	<0.0001	SFO	5.372	<0.0001
DAL	5.724	<0.0001	DAL	1.610	<0.0001
MCO	5.661	<0.0001	MCO	3.636	<0.0001
DEN	5.634	<0.0001	DEN	1.606	<0.0001
PHX	5.264	<0.0001	PHX	2.429	<0.0001
CLE	5.066	<0.0001	CLE	1.812	<0.0001
LAX	4.314	<0.0001	LAX	1.199	<0.0001
DCA	4.250	<0.0001	DCA	0.440	0.1498
SEA	3.161	<0.0001	SEA	6.195	<0.0001
MEM	2.570	<0.0001	MEM	-4.278	<0.0001
AUS	2.386	<0.0001	AUS	-0.219	0.5967
HOU	2.241	<0.0001	HOU	-0.628	0.0661
BNA	2.173	<0.0001	BNA	0.192	0.529
TPA	2.052	<0.0001	TPA	3.326	<0.0001
CMH	1.295	0.0013	CMH	0.551	0.147
OAK	-0.065	0.8586	OAK	-2.609	<0.0001
ABQ	-0.400	0.3477	ABQ	-1.794	<0.0001
SAN	-0.559	0.0911	SAN	0.401	0.2222
PDX	-0.624	0.1111	PDX	-0.058	0.8822
SMF	-1.278	0.0026	SMF	-0.129	0.7619

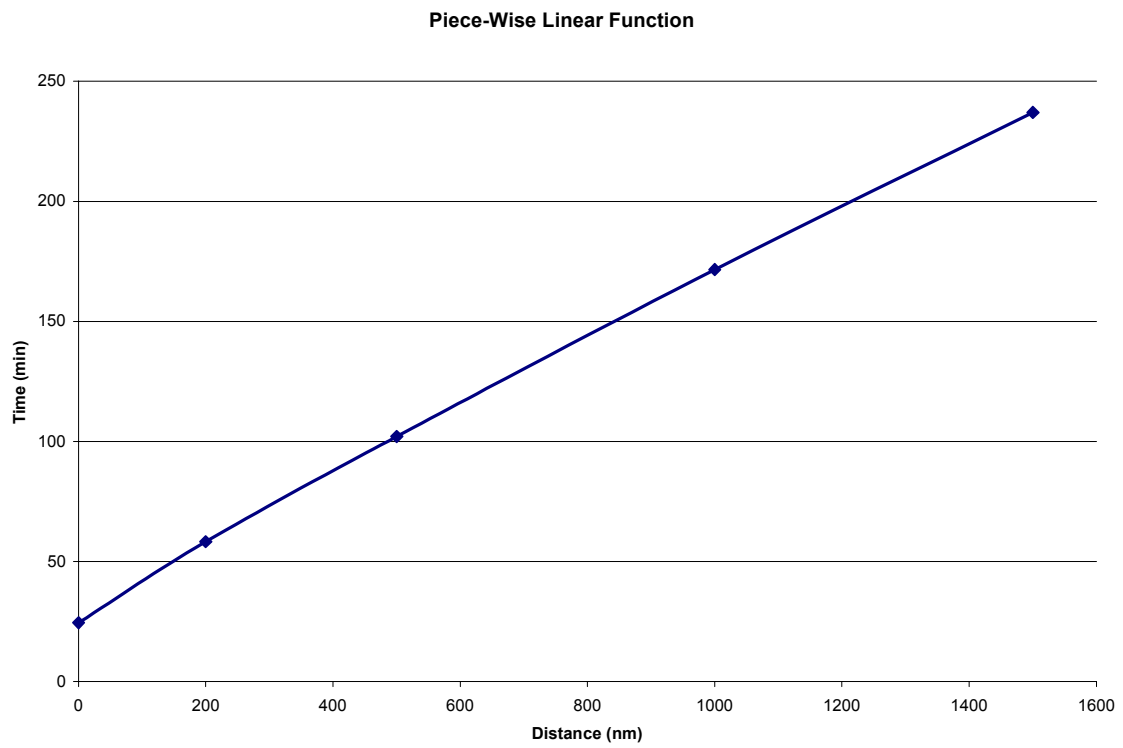


Figure 9. Piece-Wise Linear Distance-Time Function.

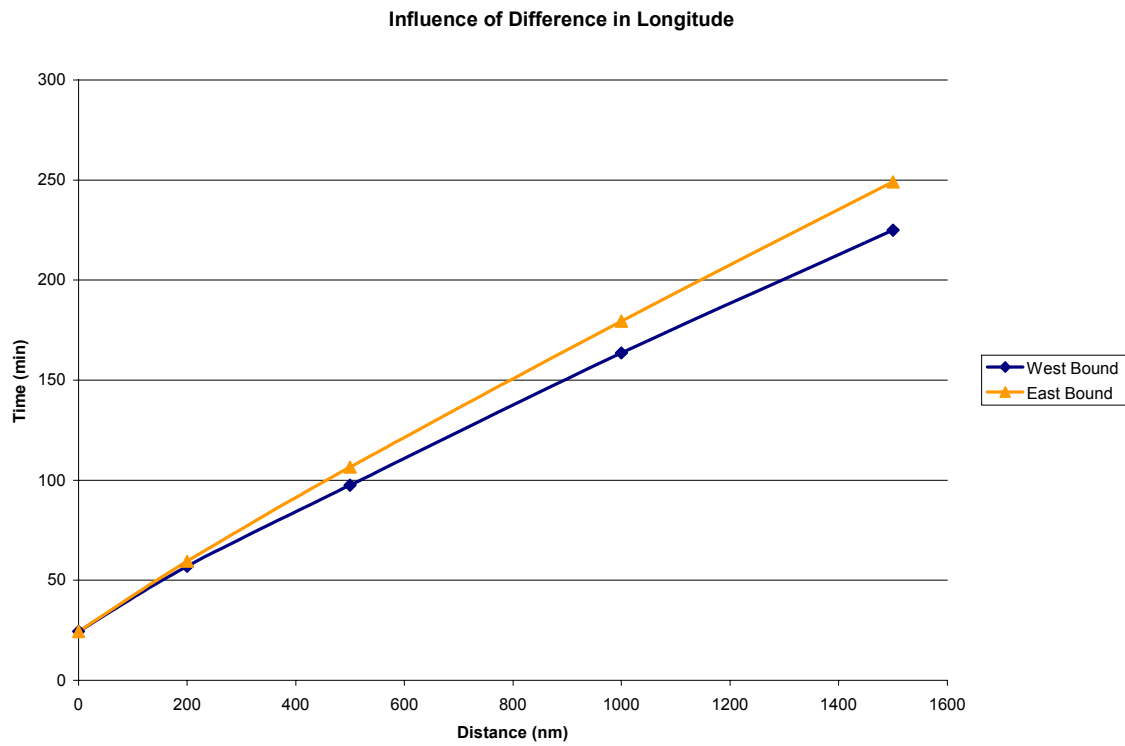


Figure 10. Influence of Difference in Longitude on East and Westbound Flights.

The URET dummy implies that flights that go through the ZID and ZME centers take about four minutes longer than those that don't. Since this effect occurs both before and after the July 2000 milestone, it is not a reflection of the impact of URET, but of other factors, such as high air traffic, that impair the operation of these centers. The AFTER variable estimate shows that the flights that took place in June 2000 had lower flight times on average than those that took place in June 1999.

Finally, and most importantly for the purposes of this study, the estimate for θ is negative. This implies that times for URET flights decreased more than times for non-URET flights between the pre- and post-milestone period. This difference, estimated here to be 3.3 minutes, can reasonably be attributed to URET itself. We will return to the discussion of this effect below.

Table 7 presents estimates of the airport fixed effects based on the June data set. We see that in the great majority of cases the fixed effects are positive and statistically significant at the 0.01 level. This is expected since fixed effects are estimated only for airports with high delays. The magnitudes of the arrival and departure fixed effects are similar, although the sources of the effects are different. For arrivals, the primary source is probably differences in the levels of congestion and delay at the various airports. For departures, the effects are most likely the result of taxi time differences arising primarily from differences in airfield geometries. (Recall that taxi-in times are not included in our flight time metric, so such differences are not included in arrival fixed effects.) The largest fixed effects for departures—all in the range of 20 minutes—are for JFK, EWR, LGA, PHL, and ORD. These same airports have the largest arrival fixed effects, which are all around 10 minutes. We re-emphasize that these are results for June only and may not be generally applicable.

6.3 URET Influence on Flight Times

The previous section discussed the detailed flight time model estimation results for a single month. We now turn to the focus of this study—the impact of URET on flight times. The first column of Table 8 presents estimates for the θ coefficient for each of the monthly data sets. As explained above, this is the coefficient that captures the effect we are looking for. It measures the average change in flight time (in minutes) that appears to result from URET implementation. The θ estimates are negative, and significant at the 0.05 level, for each month. Most of the estimates are between -1 and -2 minutes. The exceptions are June (-3.3 minutes) and July (-0.5 minutes). The small July effect may stem from the fact that two-way communication was implemented in early July of 1999, making that month, like July 2000, primarily in the after period.

It is possible to decompose the regression estimates of the flight time model into coefficients for the three flight time components: departure delay, taxi-out time, and airborne time. The regression coefficients of these component models will sum to those for the total flight time. The component estimates enable us to better understand where the flight time changes take place.

Table 8. Regression Coefficients of Individual Flight Times Model and Its Components (minutes)*

Month	Flight Time	Airborne	Departure Delay	Taxi-out	Adjusted R ²
February	-1.643	-0.239	-1.392	-0.011	0.8693
March	-1.367	-0.512	-0.929	0.074	0.8672
April	-1.354	-0.452	-0.865	-0.037	0.8591
May	-1.099	-0.196	-0.929	0.026	0.8445
June	-3.260	-0.345	-2.828	-0.086	0.8186
July	-0.502	0.223	-0.751	0.026	0.8225

*The Coefficients in bold letters are statistically significant on 1% level

Table 8 presents the component estimates. It reveals that airborne times for URET flights generally decreased about 15 seconds per flight in the in the 2000 months as compared to the 1999 months. On the other hand, much larger reductions in departure delays are observed. It appears that, somehow, URET enabled flights to depart sooner. One mechanism for this is a reduction in miles-in-trail restrictions for the URET airspace that can ultimately propagate back to the departure airport. Another possibility is that the two-way communication capability enabled controllers to more easily re-plan flights so they could avoid adverse weather, without waiting for it to clear.

The regression estimates vary considerably from month to month. The effect on total flight times is almost three times higher for June than for the other months, with the greater part of this difference involving departure delay time. A possible explanation for these month-to-month differences is that the impact of URET is weather dependent. Specifically, it may be that under adverse weather, URET enables more flexible routing and therefore a reduction in weather-related delays.

To investigate this theory, we divided the days included in our data set into categories based on average delay. Average daily delays for all days within the period analyzed were calculated from ASQP data. The distribution of this daily average is shown in Figure 11. Based on that distribution, we divided the data set into three groups: days with average delays less than or equal to 20 minutes, those with average delays between 20-40 minutes, and those with average delays greater than 40 minutes. The individual flight time model was then estimated on each set.

In this analysis, departure delay and taxi-out components are merged into one component: Time-at-Origin (TAO). From the previous results, we know that the effect of URET on TAO will be on the departure delay component. Tables 9, 10 and 11 show the estimation results, by average delay category.

Tables 9-11 suggest substantial differences in URET impacts. In general, the higher average delay, the greater the effect of URET. The greatest disparity is in the time-at-origin effect, which is generally less than a minute for low delay days, increasing to 2-3 minutes for days with moderate delays, and to 5-6 minutes on the worst days. There is some evidence of a differential impact on airborne time as well, but the difference is much smaller, and probably insignificant statistically. Also, it is important to remember when viewing these tables that in the case of July the difference is probably not attributable to URET, since two-way communication was in place for most of July of 1999.

In summary, the individual flight time analysis suggests that flight times for URET flights decreased after URET implementation, that the strongest impact was on time-at-origin, and that the impacts were greatest on days when average delays throughout the system were highest. Presumably, these impacts reflect the ability of URET flights to get into the air sooner by avoiding ground holds, particularly those associated with adverse




Figure 11. Distribution of Average Daily Delays.

Table 9. Regression Coefficients of Individual Flight Times Model and Its Components, Group 1*

Month	Flight Time	Airborne	Time at Origin	Adjusted R ²	Number of Days Used	
February	-1.365	-0.460	-0.905	0.9012	23	19
March	-0.219	-0.205	-0.013	0.8925	24	22
April	-0.409	-0.244	-0.165	0.8982	19	21
May	1.157	0.158	0.999	0.8880	18	19
June	-1.020	-0.328	-0.692	0.8841	13	8
July	-0.342	0.014	-0.355	0.8805	13	15

*The Coefficients in bold letters are statistically significant on 1% level.

Table 10. Regression Coefficients of Individual Flight Times Model and Its Components, Group 2*

Month	Flight Time	Airborne	Time at Origin	Adjusted R ²	Number of Days Used	
February	-1.068	-0.092	-0.977	0.8127	4	8
March	-4.623	-1.553	-3.070	0.8086	7	8
April	-2.189	-0.762	-1.427	0.8218	9	7
May	-4.100	-0.683	-3.417	0.8053	13	8
June	-3.023	-0.541	-2.482	0.8154	12	17
July	-1.329	0.264	-1.593	0.8145	13	10

*The Coefficients in bold letters are statistically significant on 1% level.

Table 11. Regression Coefficients of Individual Flight Times Model and Its Components, Group 3*

Month	Flight Time	Airborne	Time at Origin	Adjusted R ²	Number of Days Used	
February						
March						
April	-7.177	-0.728	-6.448	0.7659	2	2
May						
June	-5.945	-0.515	-5.430	0.7455	5	6
July	2.742	0.639	2.103	0.7540	5	6

*There was no data set with more than 40 minutes delays for February, March and May, and results in bold letters are statistically significant on 1% level.

en route weather. These results suggest that prior studies, by focusing on the portions on flights within URET sectors, overlooked a major benefit mechanism.

7. Caveats

The results from both the average flight time and individual flight time analyses are consistent in an important respect. The average flight time analysis shows improvement in average flight times of URET flights, but not in average airborne times. The individual flight time analysis shows that even though the airborne times improved slightly for URET flights in the after period, the much greater improvement was in the time at origin.

There remains, however, room for skepticism. For one thing, our analysis yields the occasional anomalous result—for example the apparent increase in flight times on moderate delay days from May 1999 to May 2000 (see Table 10). While these anomalies are clearly outweighed by the much larger number of results pointing in the other direction, they suggest the possibility that confounding factors—changes in the system between 1999 and 2000 unrelated to URET—may account for the apparent effects of URET observed here. Two of the most important of these potential confounding factors are weather and en route congestion. Here we will briefly consider these possibilities.

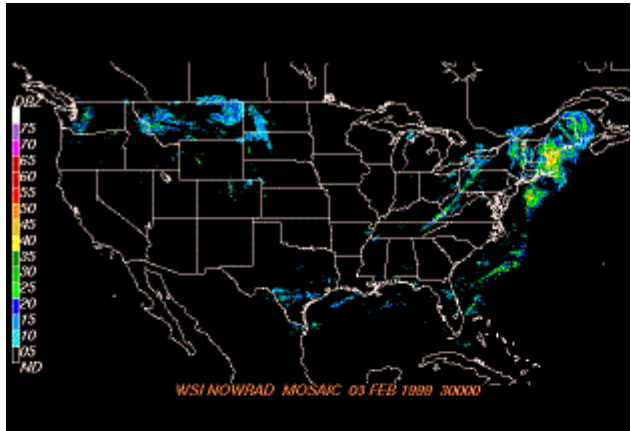
En-route airspace congestion is one reason for delays. Even without URET, flight times for flights in URET airspace might improve if the congestion in that airspace decreased. This is an especially important factor in the case of ZID, which includes some of the most congested airspace in the United States. Therefore, we compared the ZID and ZME throughputs in the before and after periods. We counted aircraft present in each center during the 11 busiest hours (13:00–23:00 GMT [10]) for each day in both periods. These counts are further averaged over the whole period and a t-test was used to discern if the differences in throughputs from before and after periods are statistically significant. Results are shown in Table 12. They show that for ZID the throughput decreased in 2000, but the difference is very small. ZME throughput practically stayed the same in both periods. This suggests that the decrease in flight times is unlikely to be the result of reduced congestion.

Weather is another potential confounding factor. Severe weather greatly influences air traffic operations. Depending on the weather situation, the NAS performance changes daily and seasonally. The analyses above deal with the weather effect in two ways. First, by comparing operations over an extended period of time individual daily weather differences may be averaged out. Second, in the individual flight time analysis, we attempted to explicitly incorporate the influence of weather through the NAS delay metric. Neither of these approaches is entirely satisfactory, however. The following example shows how weather differences could affect our results.


Consider the comparison is between two weekend days, one in February 2000 (2/13/2000) and the other in February 1999 (2/7/1999). These days were chosen for this illustration because one day had favorable weather, while the other did not. Figure 12 shows weather patterns for the two days. Figure 13 shows the AATCs obtained.

Table 12. Hourly Throughput for ZID and ZME

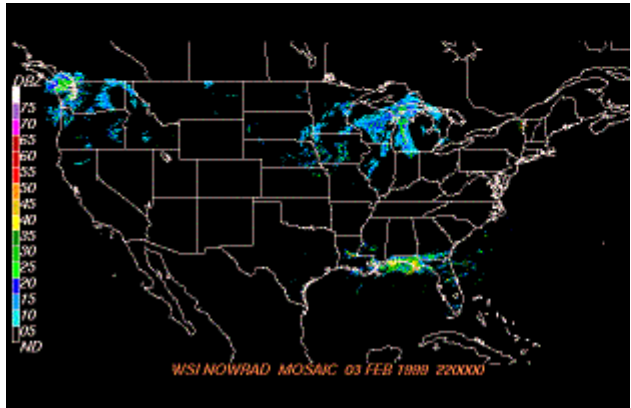
ZID			ZME		
Counts 1999	Counts 2000	P-Value	Counts 1999	Counts 2000	P-Value
397	383	0.039	322	319	0.5315



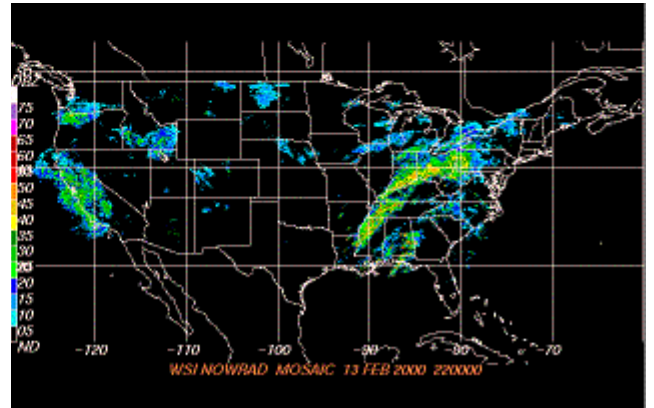
03/02/1999 0300 GMT



13/02/2002 0400 GMT



03/02/1999 2200 GMT



13/02/2002 2200 GMT

Figure 12. Radar Weather Observations for Days 2/13/2000 and 2/3/1999. (Source: <http://www4.ncdc.noaa.gov/cgi-win/wwcgi.dll?wwnexrad~images2>)

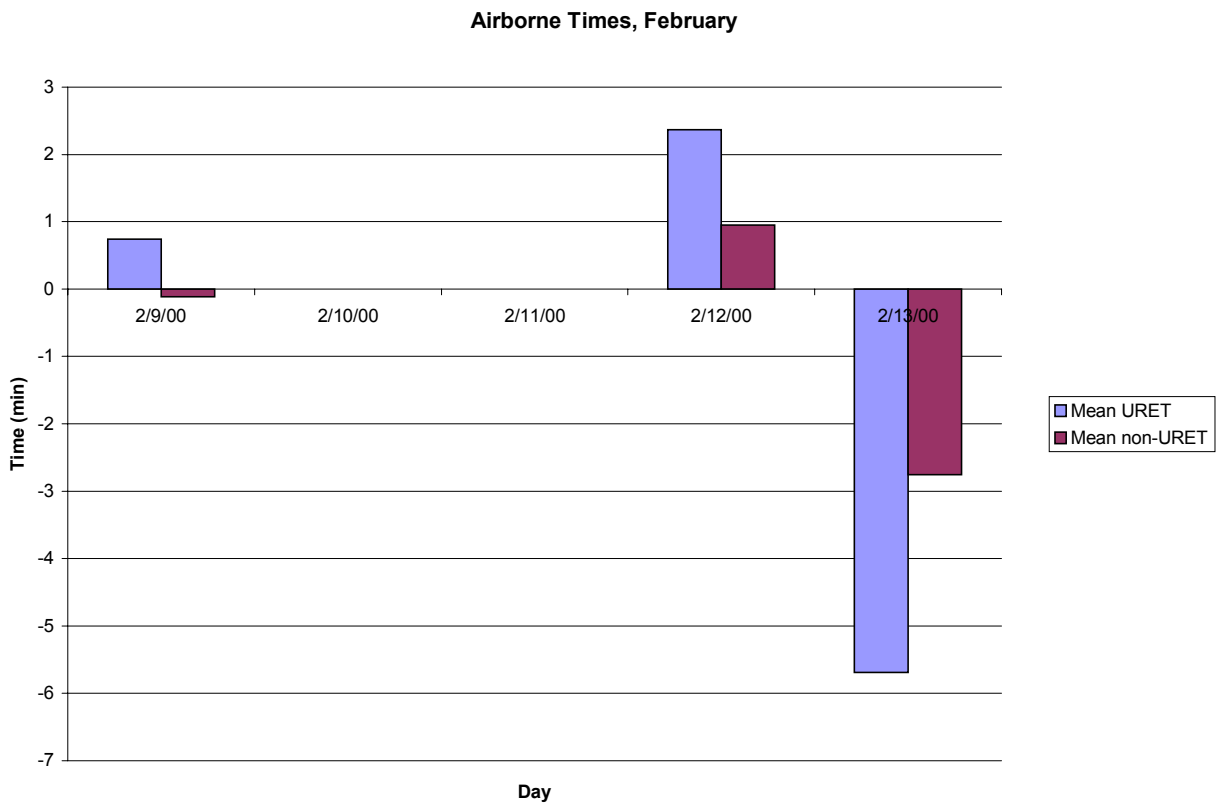


Figure 13. AATC for URET and Non-URET City-pairs.

From Figure 12 we can see that the weather was worse on 2/13/2000 than on its comparison day, 2/7/1999. Figure 13 represents the AATCs for two categories of city-pairs. Results are shown for five February weekdays, where comparison is made between years 1999 and 2000. We compared 2/13/2000 to 2/7/1999 and found that airborne flight times for URET city-pairs did not improve. In fact, URET airborne flight times increased by almost 6 minutes, while for non-URET city-pairs the increase was only 3 minutes. However, the NEXRAD images in Figure 12 show that the weather on 2/13/2000 was bad over URET centers as well as over the west coast and central USA, while on 2/7/1999 the weather was bad on the east coast but clear over URET centers. This probably explains our results, and points to the ability of weather differences to overwhelm the effects of URET, at least at the daily level.

8. Conclusions

Our analysis provides evidence that URET reduces flight times. Surprisingly, the analysis shows that the greater flight time reduction occurs not in the en route portion of the flight, but rather while the aircraft is parked at the gate. This may be because URET provides en-route controllers with a better picture of current and incoming traffic and the overall situation in sectors, so they do not have to impose miles-in-trail separations as large and as often as in other ARTCCs. It may also derive from expedited clearance procedures that enable controllers to respond to more nimbly to adverse weather conditions.

Even though the greater part of the impact is on the time-at-origin, the individual flight times analysis suggests that URET reduced airborne time by an average of 15-30 seconds per flight. It is instructive to compare this with work by other researchers. The FFP1 Metrics Team report [10], reported a Distance Saved metric, which is the sum of the reductions in lateral distance flown through the URET centers as result of flight plan amendments. The Distance Saved values for ZID in May and June of 1999 were approximately 200 and 500 miles respectively; in May and June of 2000 they were 4600 and 3500 miles. Based on the number of flights through ZID in this month, this translates into an average distance savings of 0.9 and 0.6 miles per flight respectively. Assuming a cruise speed of 7 miles/min, this yields time savings per flight is 7 and 5 seconds for the two months. Our results thus suggest airborne flight time savings that are at least double those derived from flight plan amendments within the URET centers. The difference could reflect airborne time savings that accrue outside the URET airspace, either because of additional distance savings or reduced miles-in-trail restrictions. Wind effects may also play a role.

The analysis methodology used in this report is a promising beginning for a posteriori benefits assessment of air traffic management and air traffic control innovations. The methods are global in that they capture benefits throughout the flight, and can be translated easily into monetary terms because they consider flight times. The individual flight time analysis has the added advantage of exploiting the vast, and growing, quantities of data about NAS operations that are becoming available. Future work will take this concept further by adding effects related to en route weather, time-of-day, and other factors. It will also investigate how URET flight time impacts vary over time and for different flight lengths. A further enhancement is to link quantitative estimation results with a qualitative view of the benefit mechanisms that URET engenders. Our results suggest that present understanding of how URET generates benefits may be very incomplete, and resulting in vast underestimates of URET benefits. This research is a first step in remedying this situation, but future work is necessary to verify our findings, and translate them into benefits estimates of investment-grade quality.

References:

1. W. C. Arthur, M. P. McLaughlin, **URET Interfacility Conflict Probe Performance Assessment**, MITRE MP 98W0000204, December 1998
2. K. Kerns, A. L. McFarland, **Conflict Probe Operational Evaluation and Benefits Assessment**, MITRE, MP 98W0000239, November 1998
3. D. J. Brudnicki, A. L. McFarland, **URET Conflict Probe Performance and Benefits Assessment**, MITRE, MP 97W0000112, June 1997
4. Daniel J. Brudnicki, Alvin L. McFarland, Susan M. Schultheis, **Conflict Probe Benefits to Controllers and Users Indications from Field Evaluations**, MITRE, MP 96W0000194, August 1996
5. Joseph C. Celio et al, **Free Flight Phase 1 Conflict Probe Operational Description**, MITRE, March 2000
6. **Benefits for NAS Users Enabled by URET: A Strategic Planning Tool**, FAA, MITRE, 4th USA/Europe ATM R&D Seminar
7. Joseph C. Celio, Marilyn G. Walker, Alvin L. McFarland, **User Request Evaluation Tool (URET) Benefits During Free Flight Phase 1**, MITRE, July 2000
8. Michael J. Burski, Joseph Celio, **Restriction Relaxation Experiments Enabled by URET a Strategic Planning Tool**, MITRE, June 2000
9. Daniel B. Kirk, Winfield S. Heagy, Alvin L. McFarland, Michael J. Yablonski, **Preliminary Observations About Providing Problem Resolution Advisories to Air Traffic Controllers**, June 2000
10. Free Flight Phase 1 Metrics Team, **FFP1 Performance Metrics Results to Date**, December 2000, June 2001 (ffp1.faa.gov)
11. **Aviation Capacity Enhancement CD-ROM and Airport Database**, 1998.

List of Abbreviations:

AATC	Average Airborne Time Change
AERA	Automated En-Route Traffic Control
AFTC	Average Flight Time Change
APD	Automated Problem Detection
ARTCC	Air Route Traffic Control Center
ASQP	Airline Service Quality Performance database
AT	Airborne Time
ATC	Air Traffic Control
CAASD	Center for Advanced Aviation System Development
CDM	Collaborative Decision Making
CHI	Computer Human Interface
FAA	Federal Aviation Authority
FFP1	Free Flight Phase 1
FFP2	Free Flight Phase 2
FT	Flight Time
ETMS	Enhanced Traffic Management System
HCS	Host Computer System
NADIN	National Airspace Data Exchange Network
NAS	National Airspace System
nm	Nautical Mile
PARR	Problem Analysis Resolution and Ranking
TAO	Time at Origin
URET	User Request Evaluation Tool
US	United States
ZID	Indianapolis ARTCC
ZME	Memphis ARTCC