NRA Task Order 47

Refined Benefits Assessment of the Collaborative Arrival Planner (CAP) AATT Decision Support Tool

Prepared By
Metron, Inc.

Author: Robert Hoffman

NASA Technical Monitor: Chris Scofield
Metron Principal Investigator: Robert Hoffman

Task Team:
- Chris Brinton
- Mike Feldman
- Robert Hoffman
- Kevin Kollmann
- Dejan Neskovic
- Gretchen Wilmouth

Contract with NASA Ames Research Center, under subcontract with Titan Systems Corporation/System Resources Division.

November 17, 2000

Executive Summary

CAP is a proposed interface between the User community and Air Traffic Service Providers at airports in which CTAS has been deployed. This study performs a benefits assessment of the CAP tool as part of the Advanced Air Transportation Technologies (AATT) Project within NASA's Aviation System Capacity (ASC) Program. The ultimate use of the study will be to compare CAP with other proposed AATT tools. This study is distinguished from prior CAP studies by the absence of proprietary data, thus making it available for distribution to other AATT contractors.

CAP was in a conceptual stage of development when this study was performed. Hence, it was necessary to make a number of assumptions regarding CAP design. We assumed the existence of the following functions within CAP.

- *Passive information sharing* that would display to air carriers the same CTAS/TMA advisories that are displayed to controllers in the Terminal Radar Approach Control (TRACON) and Air Route Traffic Control Center (ARTCC). This will tell the air carriers what arrival times the controllers are working toward and which runways will be used by their aircraft.
- *Acceptance of air carrier preferences* for arrival priorities during periods of delay due to traffic or weather induced congestion. This function is currently performed ad hoc via telephone calls from the air carriers to the ATSP. We assume that preferences would be expressed on a flight-by-flight basis using simple flight delay cost structures.
- *An algorithm to equitably balance preferences* from multiple Users with possibly competing interests. One possible such algorithm was developed by Metron in a CAP Equity study.
- *Real-time access to the controlled times of arrival* (CTAs) set during a ground delay program (GDP). CAP would work to restore flights to these CTAs, thus preserving the equitable distribution of constrained arrival capacity issued during a GDP.

A blanket assumption of this study is that any advisories posted by TMA or the CAP tool for arrival time adjustments will be carried out by ASTP. For this reason, many of our results are on the optimistic side.

Next, we give a brief description of each of the anticipated benefits of the CAP tool, the metric used to evaluate it, the corresponding AATT high-level benefits category, and our final assessment of the benefit.

1. Improved Aircraft Arrival Prediction Accuracy (Section 4.2)
This would result from using CTAS-generated ETAs rather than the ETMS-generated ETAs (Section 5.1). AATT Benefits Category: *Predictability.* For each arriving aircraft at select airports, and for each ETA-generating system (CTAS and ETMS), the Integrated Predictive Error (IPE) metric was used to form a weighted average of ETA errors (ETA – actual arrival time) captured at sequential times during the last 60 minutes of flight. This generated on IPE value for each flight. Statistical analyses were then performed.

We found that CTAS-generated ETA streams are on average between 30.28 seconds (at DFW) and 64.00 seconds (at DEN) more accurate in absolute value than ETMS-generated ETA streams. Both systems, CTAS and ETMS, are biased toward the positive, meaning that flights tend to arrive earlier than predicted. For instance, on 9/19/00 at DEN, the average CTAS ETA error was175 seconds; the average for ETMS was 387 seconds.

2. Improved Basis for Making Hold/Go and Diversion Decisions (Sections 4.3, 4.4)
Access to improved ETAs should allow air carriers to make more informed decisions whether to hold outbound aircraft for connection with potentially delayed inbound aircraft at a hub airport. Also, this should allow them to make a more informed decision whether to divert a flight from its original destination. AATT Benefits Category: *Predictability*. The metric used was the number of flight instances in which exactly one of the two systems (CTAS or ETMS) generated an ETA that fell within a specified range of the actual time of arrival of the flight and the other system fell outside the specified range.

We found that out of 8,573 flight instances captured at three airports over at total of 16 days, there were only eight instances in which one of the two systems prevailed over the other. Thus, we were unable to determine whether CTAS would provide an improved basis for making Hold/Go or diversion decisions.

3. Airborne Holding and Expected Clearance Times (Section 4.5)
In times of heavy demand or constrained arrival capacity, CTAS provides to air traffic service providers a recommended schedule of arrival times that will make efficient use of airport resources. We anticipate that access to this schedule would provide air carriers with estimates of release time from airborne holding that are more accurate than currently available to them. AATT Benefits Category: *Predictability*. We used the IPE statistics but only for those flights subject to airborne holding.

We found CTAS a slightly more accurate predictor of arrival times than ETMS. We anticipated that this performance gap would be significantly wider for flights subject to airborne holding. However, this requires a list of flights subject to airborne holding during our data collection period. We were not able to obtain (nor model) this flight list. Moreover, given the fair weather conditions present at all airports during our data collection period, we believe that, even if such a flight holding list were available, there would not be sufficient flight holding instances to reach a solid conclusion.

4. Savings from an *Inter*-carrier Delay Exchange Program (Section 5.3)
Simulations were conducted in which flights were processed in order of increasing ETA, to simulate their entrance to the CTAS domain of operation. AATT Benefits Category: *Flexibility*. Flights were assigned CTAS-scheduled arrival times according to demand, airport capacity, and minimum separation. Using flight cost structures as a guide, delay exchanges were made between flights to move flights into lower delay buckets.

We estimate that for the year 1999, there would be an improvement in on-time arrival performance of 1.81%. NAS-wide, which corresponds to 152,931 more on-time arrivals. At $49.90 per delay minute, this amounts to $29,914,527 savings over the course of a year.

5. Savings from an *Intra*-carrier Delay Exchange Program (Section 5.4)
AATT Benefits Category: *Flexibility*. We examined the results of the inter-carrier delay exchanges of Section 5.3 for benefits derived when a carrier is restricted to trades with flights of the same carrier. We found that there would be a 2.01% increase in on-time arrival performance, which corresponds to 2,855 more on-time arrivals per year for a single hubbing operation, for annual savings of $558,339.

6. Savings from Arrival Time Adjustments toward Carrier Schedules (Section 5.5)
In the inter-carrier delay exchange simulation, flights were moved into arrival slots occupied by other flights. In this analysis, we simulated the movement of flights into vacant arrival slots. Flights were sped up, when feasible, to get them as close as possible to their carrier-scheduled arrival times. AATT Benefits Category: *Flexibility*.

We estimate that there would be a 2.06% increase in on-time arrival performance. This amounts to 174,054 more on-time arrivals across the NAS per year, for annual savings totaling $21,740,090.

7. Restoration of GDP-issued Controlled Times of Arrival (Section 5.6)
Upon execution of a GDP, flights often arrive outside of their designated arrival time windows for various reasons. Two analyses were performed to assess the potential for incorporation of airline objectives into terminal area scheduling via the integration of CTAS and the Collaborative Decision Making (CDM) program. AATT Benefits Category: *Access*. The metric used was percentage of flights that could be restored to their CTAs in a GDP.

Based on the averages over all GDP test days at PHL, we found that the percentage of flights at PHL that could be restored to their controlled time of arrival (CTA) by CTAS lies somewhere between 44.73% and 21.98% (40.27% and 19.21% at SFO), with the most likely estimate lying closer to the latter figure(s).

Summary Tables for the Study: Table 6-1 and Table 6-2 in Section 6 of the body of the report provide a summary of metrics and benefits for the years 1999 and 2015, respectively.

1 Introduction

This study develops and applies a methodology for the assessment of the benefits that would arise from deployment of the Collaborative Arrival Planning (CAP) decision support tool of the Advanced Air Transportation Technologies (AATT) Project, as part of NASA's Aviation System Capacity (ASC) Program.

CAP is designed to aid the Airline Operations Centers (AOCs) and air traffic service providers (ATSPs) in the planning of terminal-area scheduling of aircraft arrivals. Currently, control of aircraft within the terminal area is set by ASTPs with little input from the AOCs. CAP would allow the AOCs to interact with this process. The information-sharing aspect of CAP would provide to the AOCs real-time information on aircraft movements and timing within or close to the terminal area. Air carrier operations are highly dynamic operations that require constant manipulation and awareness of any delays that might jeopardize gate management, crew and equipment connections, and passenger connections. This flow of information toward the AOCs would allow them to mitigate potential damages associated with delayed arrivals.

In addition to information-sharing, CAP would allow air carriers to input arrival priority requests for their own flights to the terminal area scheduling process during periods of delay due to either traffic or weather induced congestion for arrivals. This would greatly enhance the current state of communications in which telephone calls are made from the AOCs to the Air Traffic Service Providers (ATSP). Subsequent versions of CAP may supplement CTAS data with ETMS data to aid in the planning of flights from origin to destination, rather than just at destination. This may follow a paradigm similar to those established by the collaborative decision making (CDM) program.

In Fiscal Year 1998, the Advanced Air Transportation Technologies (AATT) Project funded an initial CAP benefits assessment, "Benefits Assessment of Airline-influenced Arrival Scheduling", which explored the economic benefits derived from manipulation of arrival sequences at hub airports during peak traffic periods. In Fiscal Year 1999 (FY99), another CAP benefits assessment was funded. However, these studies contain much proprietary information and, therefore, are not publicly available. There is a need for benefits assessment that could be made available to the entire AATT team and that would feed larger-scale benefits assessments made by AATT.

This study is distinguished from prior studies on CAP benefits in that it contains no proprietary data. This allows for dissemination to and reference by an audience beyond the government. Specifically, it can be distributed to other AATT contractors and feed larger AATT assessments. The set of AATT high-level categories of benefits have been used as a framework, though not each of these has a benefit associated with it. There are additional categories of benefits uncovered in this development phase that are not readily attributable to the AATT high-level categories.

This document is divided into six sections. Section 2 provides a description of the capabilities and functions of CAP that we have assumed for the purposes of conducting the study. Section 3 describes the nature and source of data used in the study. Sections 4 and 5, the core of the document, cover the assessment of CAP benefits. These benefits naturally group themselves into two categories: those associated with the output of CAP that would be provided to air carriers (Section 4) and those associated with preferences that would be provided by the air carriers to the CAP tool (Section 5). This sectioning is largely done for ease of exposition; analyses of a similar nature share a dependency upon technical explanations, analytic techniques, and CAP functionality. Each subsection is clearly labeled with the AATT benefits category under which it falls.

Table 6-1 and Table 6-2 in Section 6 of the body of the report provide a summary of metrics and benefits for the years 1999 and 2015, respectively.

2 Assumptions Regarding CAP Functions

CAP can be viewed as an interface between CTAS operations and the User community. Since CTAS is airport specific, we assume that each CTAS installation will have associated with it its own, possibly unique, version of CAP. For the purposes of benefits assessment, we also assume that CAP will reside at each site in which CTAS currently resides.

At the time of this study, it is uncertain exactly what the decision support capabilities of the Collaborative Arrival Planner (CAP) will be. However, based on dialogue with the CAP development team and based on prior research efforts into CAP (References [1], [2], [7], [8], [9]), we have identified a set of potential CAP features that are both realistic and promising. Those features can roughly be divided into three categories:

1. *CTAS to AOC Data Transmission*: This is a mechanism to transmit and display at the air carrier operational centers (AOCs) estimated time of arrivals (ETAs) of flights inbound to the CTAS airport. This would allow the AOCs to have access to CTAS ETAs and CTAS scheduled slot times, as displayed in TMA timelines.

2. *User Input to CTAS Scheduling*: This is the infrastructure necessary for NAS Users to affect arrival scheduling performed by CTAS TMA software. Strictly speaking, CTAS does not control aircraft. Rather, it provides recommendations to air traffic service providers who perform the actual control. An essential assumption in this study is that all recommendations made by CTAS can be carried out by air traffic service providers. Section 5 discusses in more detail the assumptions of the form that this input would take.

3. *A CAP-FSM Interface*: The Flight Schedule Monitor (FSM) is the primary tool by which air traffic service providers monitor arrival demand at U.S. airports and by which they implement ground delay programs. These ground delay programs issue controlled arrival times (CTAs) to flights bound for a capacity impaired or demand inundated airport. These CTAs are set by AOCs and by the FAA to maintain equitable assignment of scarce airport resources. It has been envisioned that an interface between CTAS and FSM would provide the CTAS TMA tool with knowledge of the CTAs. TMA could then work toward restoring flights to their CTAs when they have drifted off them.

Feature 1 will provide benefits mainly in the form of increased Predictability. Feature 2 provides benefits mainly in the form of Flexibility, Access, and Equity. Feature 3 will provide benefits in the form of Equity. These benefits are discussed in more detail in the respective sections of the document.

3 Data and Airport Selection

Several of the analyses conducted in this study are centered on User inputs to CTAS scheduling of arrivals at a CTAS airport. Since one of the primary objectives is to provide a document that can be made generally available, we have avoided use of or reference to AOC proprietary data. Instead, we have relied exclusively on two sources of data: CTAS and ETMS. Since CTAS-oriented data is much harder to obtain than ETMS data, this has naturally placed a limit on the range of analyses that can be conducted.

3.1 CTAS Data

We collected System Analysis Recording (SAR) data for the following airports and respective dates:

- DEN 9/14/2000 through 9/20/2000
- DFW 8/20/2000 through 8/26/2000
- PHL 9/5/2000 through 9/6/2000

The SAR data was converted to CTAS data by running it through Metron's version of the CTAS TMA software. The CTAS data fields that were generated/collected were

- Runway Arrival Time
- TMA Scheduled Slot Time at *T* minutes prior to arrival where *T* = 5, 10, 15, …, 60.
- Estimated Time of Arrival at *T* minutes prior to arrival where *T* = 5, 10, 15, …, 60.

This simulated a complete history for each flight over the last hour if its flight.

3.2 ETMS Data

The Enhanced Traffic Management System (ETMS) data was extracted from *Delta* files, which are generated by the Flight Schedule Monitor (FSM), based on the aggregate demand lists (ADLs) that Volpe distributes. An ADL is a chronological sequence of flight record data blocks. FSM condenses the ADL file into a Delta file by storing a flight record only if one of its fields differs from the field in the corresponding flight record in the prior data block.

Several of our metrics relied on CTAS versus ETMS comparisons. Since the SAR data necessary to generate CTAS data was limited to one week for DEN and DFW and to two days for PHL, this limited the days and airports for which comparisons could be made.

Airport Selection Criteria

The driving factor in selecting airports for the study is the availability of data. ETMS data is generally available, but CTAS data is harder to obtain or generate. Although the data to be fed into CTAS simulators is also available, in order to run the simulator, the TMA software must be adapted to the airport, meaning that it must be adapted to highly airport specific settings such as the runway configurations, meter fix locations, viable approach paths, etc. At the time of this study, CTAS has been adapted to only the following airports: DFW, PHL, DEN, MSP, LAX, ATL, MIA. Of these, we have selected DFW, PHL and DEN to represent a cross section of the CTAS airports.

Since Section 5 is not dependent only upon ETMS data, we were able to include ORD as one of our airports of study. ETMS data in this section covered the first six months of 1999.

4 Benefits Assessments Related to ETA Accuracy

4.1 The Integrated Predictive Error (IPE) Metric

Several of the metrics employed in this section required an assessment of the quality of predictions made for arrival times. Since these predictions are made at varying times in the path of a flight, we used The Integrated Predictive Error (IPE) metric evaluate overall quality of the streams. The metric was originally developed for analysis of long-term trends in departure prediction accuracy (see [3] and [4]) but it applies equally well to arrival predictions. Section 4.1.1 explains the workings of the IPE metric. If desired, the reader may wish to skip this section with the understanding that IPE is a weighted average of absolute predictive errors made over time for a single event.

4.1.1 Explanation of the IPE Metric

There have been a number of analyses of CTAS ETA accuracy over ETMS ETA accuracy (e.g., see Reference [1]). These analyses generally use the ‘snapshot’ method in which a time *T* minutes prior to arrival is fixed for all flights in the study. The accuracy of the CTAS ETAs and the ETMS ETAs is compared by examining the distribution of the CTAS ETA errors and the ETMS errors. Error is determined by

ETA_error = ETA – ARTA (or ETA – MFTA),

Where

ARTA = actual runway time of arrival
MFTA = meter fix time of arrival.

There are two problems associated with the snapshot method. First, a distribution must be formed for each value of *T* that is selected. For instance, if *T* is set to 5, 10, and 15 minutes, then three distributions are generated. The IPE metric avoids generating multiple the need to evaluate the two systems, CTAS and ETMS, at each value of *T*. Instead, one comparison is made over all values of *T*.

The other problem with the snapshot method is that the results may vary with the values of *T* chosen. It is possible (though not likely) that one ETA-generating system could prevail over each of the chosen *T* values but be worse at some of the neglected *T* values. IPE avoids this by computing a weighted average of all ETAs made for a given flight.

The snapshot method is perfectly appropriate for analyses in which *T* can be set to capture an essential point in the path of a flight. For instance, in Reference [2], *T* was set to 30 minutes before arrival. This corresponds to the time at which American Airlines flight crews typically provide an arrival time update once they receive the destination airport’s Automated Terminal Information Service (ATIS). In the present study, we employ the snapshot method when similar needs arise but for the overall evaluation of predictive streams, we rely on IPE to produce a single value for each flight. This is ideal for aggregation over numerous flights.

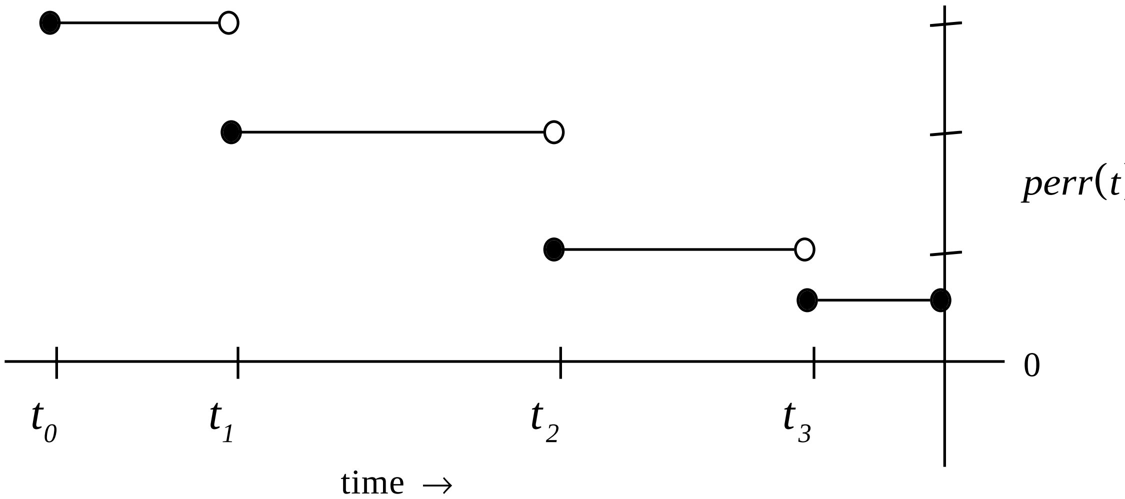

Figure 4-1: Integrated Step function

4.1.2 Sample IPE Computation

Suppose that a system has submitted an estimated time of departure ETA_t for a flight at each of the times t_0, t_1, t_2, and t_3,. Let *ARTA* be the actual time of departure of flight *f*. Then the predictive error at time *t* is defined by $perr(t) = |ETA(\text{t}) - ARTA|$. Under the assumption that the most recent update submitted by the system is its current best estimate, we can represent these updates by line segments, as in Figure 4-1. The vertical axis of the graph intersects the horizontal axis at the point in time at which the flight actually arrived. The black dots correspond to the errors of the respective ETA's. A high dot indicates a poor estimation. A zero value indicates a perfect prediction.

By going uncorrected for a length of time, each ETA defines a block of area under each of the steps in Figure 4-1. By integrating the step function over the entire time period, the IPE metric sums these areas to arrive at a final IPE value for the flight. An estimate left uncorrected for a long period of time will add significantly to the value of the IPE. Therefore, lower IPE values reflect a better stream of predictions. The lowest possible value is of course zero, indicating that there was no error in any of the predictions.

4.1.3 Units of IPE and Normalization

Since an IPE value is the result of integrating predictive errors over time, the units of IPE are time-squared. For instance, if both the predictive error and the interval of integration are measured in minutes, then IPE is in units of minute-minutes (or minutes squared).

When IPE is applied to the predictive accuracy of the arrival time of a flight, a more practical unit can be formed by measuring the predictive error in minutes and the time span in hours. The final units are minute-hours which have the following natural interpretation: suppose that the stream of predictions (ETA's) is measured over one hour.

Then an IPE of 10 minute-hours is equivalent to an error of exactly 10 minutes for the entire one-hour period. Phrased another way, the average error over the one-hour period is 10 minutes.

When comparing the relative quality of two predictive streams tracked over time horizons of different lengths, then the metric must be normalized. For instance, suppose that each of two systems reported a constant 10-minute error in arrival prediction for the same flight. Each would have a raw (unnormalized) IPE value of 10 minute-hours. However, if one of the systems were tracked over 2 hours, then it would have a raw IPE value of 10 + 10 = 20 minute-hours, thus giving the illusion that it yielded twice as much error as the other system. Normalization is achieved by dividing each raw IPE value by the length of time that it was tracked. In our hypothetical example, this would produce two normalized values:

System 1: 10 min-hours / 1 hour = 10 minutes
System 2: 20 min-hours / 2 hours = 10 minutes

This normalization also makes for easy interpretation of the metric. For eaqch of the two systems, the "average" arrival prediction error was 10 minutes. The notation that is used is as follows. Let I be the raw IPE value (integration of the step function). Let n be the length of time over which the errors were tracked (integrated). Then "ipe_n" denotes a normalized IPE value computed over track length *n*. For instance, ipe_6 is an IPE value computed over 6 hours. Henceforth, we will consider all IPE values normalized.

Note: For the purposes of this analysis, the IPE values were normalized mainly for ease of interpretation. The time periods over which the ETA streams were generated for each of the two ETA-generating systems were taken to be the same.

We will be using the IPE metric to generate a single value for the entire stream of predictions for a given flight. For instance, for each flight, there will an IPE value for the CTAS stream of predictions and an IPE value for the ETMS stream of predictions. The lower IPE value is the one with better performance.

We note that the use of the IPE metric does not preclude analysis of errors at particular points in time. For instance, one can use ipe1, ipe2, …, ipe6.

4.1.4 Variations on IPE

When computing an IPE integral, it may be desirable to devalue the error of estimates that were formed long before the actual event. For instance, in the case of departure prediction, an error of 60 minutes made 12 hours prior to arrival does not seem as significant as an error of 60 minutes made 30 minutes before arrival. One variation on the IPE metric is to apply a dampening function to the step function before it is integrated. See [5] or [6] for details. We did not see a need to dampen predictive errors in this study.

4.2 Aircraft Arrival Prediction: ETA Accuracy

Metric: Integrated predictive Error (IPE)
Data: CTAS-simulated ETAs, ETMS-generated ETAs.
AATT Category: Predictability

One of the anticipated benefits of CAP is to provide CTAS data to the AOCs. In particular, air carriers should benefit from improved accuracy in estimated times of arrival (ETA) for their aircraft when in the later stages of flight. There are two reasons to expect improved ETAs from CTAS. First, CTAS receives benefit of radar updates every 12 seconds for each flight, whereas ETMS receives benefit of radar updates every 60 seconds. Second, the trajectory modeler in CTAS is generally more refined than the trajectory modeler in ETMS. The latter factor is less of a contributor to improved CTAS accuracy. Any improvements that are gained by CTAS in arrival predictions can be decomposed into these two factors. However, since we are only concerned with the overall gain in arrival prediction accuracy, we will not be decomposing ETA predictions into these two levels. Instead, we will assess the net gain achieved by CTAS.

In the CTAS data that we have collected, each flight has a CTAS-generated ETA captured T minutes prior to arrival, for each of the values T = 60, 55, …, 5 minutes. This generates a 'stream' of arrival predictions. (Note: short-haul flights of flying time less than 60 minutes may not have an ETA for each of these values.) From ETMS-based Delta files (see Section 3), we also have ETMS-generated ETAs for approximately the same time intervals. This generates an ETMS stream of arrival predictions for each flight.

Rather than compare these ETAs on an individual basis, we have used the IPE metric to compute a single value for each of the streams. For a given stream of ETAs, the IPE value is a weighted average of the absolute errors of each ETA. The ETA error is defined as

$$\text{ETA\_error} = \text{ETA} - \text{ARTA}$$

where ARTA is the actual runway time of arrival. ARTA was taken to be the CTAS recorded arrival time (ETMS-generated arrival times are slightly less accurate).

Our evaluation was restricted to flights common to both our CTAS database and to our ETMS database. For each of these flights, we generated an IPE value for the CTAS predictive stream and one for the ETMS predictive stream. Since the CTAS streams were limited to the minimum of the last 60 minutes of flight and the flying time of the flight, we limited the ETMS to the same time interval.

Table 4-1, Table 4-2, and Table 4-3 give the compiled results for the three CTAS airports of our study, DEN, DFW and PHL, respectively. For each airport/date, a distribution of IPE values was formed, where there is one IPE value for each flight. Since the IPE metric records absolute errors, lower values indicate better performance. Each distribution is represented by its mean value, maximum, minimum, and standard deviation. Values are given in seconds. To understand the interpretation of the values, consider the 331 second value in the upper left corner of Table 4-1. This states that over all flights arriving at DEN on 9/14/00, the mean (absolute) error in CTAS arrival prediction was 331 seconds, or, about 5.52 minutes (CTAS values are shaded). Note that to the right of this value, the corresponding mean error for the ETMS system (337 seconds) is 40 seconds higher.

Scanning across shaded and unshaded column pairs (CTAS and ETMS) reveals that on each of the seven test days at DEN, the mean CTAS error was less than the mean ETMS error. The differences of the means were 40, 35, 65, 83, 97, 91, and 36 seconds, respectively, for a mean of 64 seconds over all seven days. In other words, CTAS is on average slightly more than one minute more accurate a predictor of arrival time at DEN than ETMS.

Comparable results were found at DFW and PHL (Table 4-2 and Table 4-3). CTAS outperformed ETMS on each of the seven day and averaged 32 seconds less error than ETMS at DFW, and 99.5 seconds less error at PHL.

An interesting observation is that, although the CTAS mean predictive error is less than ETMS mean predictive error, CTAS generally has higher maximum values and higher standard deviations, indicating a greater spread in values. Statistically, this can be explained by examining the distribution of IPE values for a single day. Figure 4-2 is a frequency distribution of ETA errors (IPE values) for both CTAS and ETMS at DFW on 8/21/00. The CTAS distribution has a longer 'tail', which accounts for it having a higher maximum than the ETMS distribution. This also accounts for the slightly higher standard deviation. However, the important feature of this Figure 4-2 is that the CTAS errors are bulked more toward zero.

One drawback of using the IPE metric is that it does not tell us whether the distributions of predictive errors for the two systems are biased toward the positive or negative. To address this question, we formed distributions of ETA errors at 60 minutes out before arrival. Figure 4-3 shows the distribution of errors (without having taken absolute value) at DEN on 9/19/00. This is typical of the results we found at all three airports. Both the CTAS and the ETMS distribution are positive biased, with means of 175 seconds and 387 seconds, respectively. That is, on average, flights arrive 2.9 minutes earlier than CTAS predicts and 6.45 minutes earlier than ETMS predicts. The CTAS distribution is wider and flatter than the ETMS distribution.

Since it is difficult to evaluate the impact that provision of improved arrival time estimations for individual flights would have an air carrier operations, we have not attempted to translate the benefits shown in this section to other units (e.g., dollar amounts).

DEN Average IPE Values, CTAS vs. ETMS														
	9/14/00		9/15/00		9/16/00		9/17/00		9/18/00		9/19/00		9/20/00	
Avg	331	371	349	384	322	387	306	389	299	396	277	367	315	351
Std Dev	134	160	132	146	136	149	159	170	149	146	152	149	204	175
Max	919	1199	1580	871	990	1185	1314	1357	922	979	1006	792	1706	1051
Min	48	33	46	52	33	100	35	70	14	53	37	33	4	5

Table 4-1

DFW Average IPE Values, CTAS vs. ETMS														
	8/20/00		8/21/00		8/22/00		8/23/00		8/24/00		8/25/00		8/26/00	
Avg	146	200	168	189	161	200	172	192	170	199	170	195	163	197
Std Dev	97	96	106	101	102	100	104	98	109	97	110	106	104	98
Max	876	560	789	615	707	604	864	653	1189	704	856	642	827	641
Min	17	13	13	6	3	21	3	3	18	19	8	9	13	10

Table 4-2

PHL Average IPE Values, CTAS vs. ETMS														
	9/5/00		9/6/00											
Avg	261	394	276	341										
Std Dev	149	174	214	168										
Max	1154	843	1039	921										
Min	54	26	44	41										

Table 4-3

CTAS values (shaded)

ETMS values (unshaded)

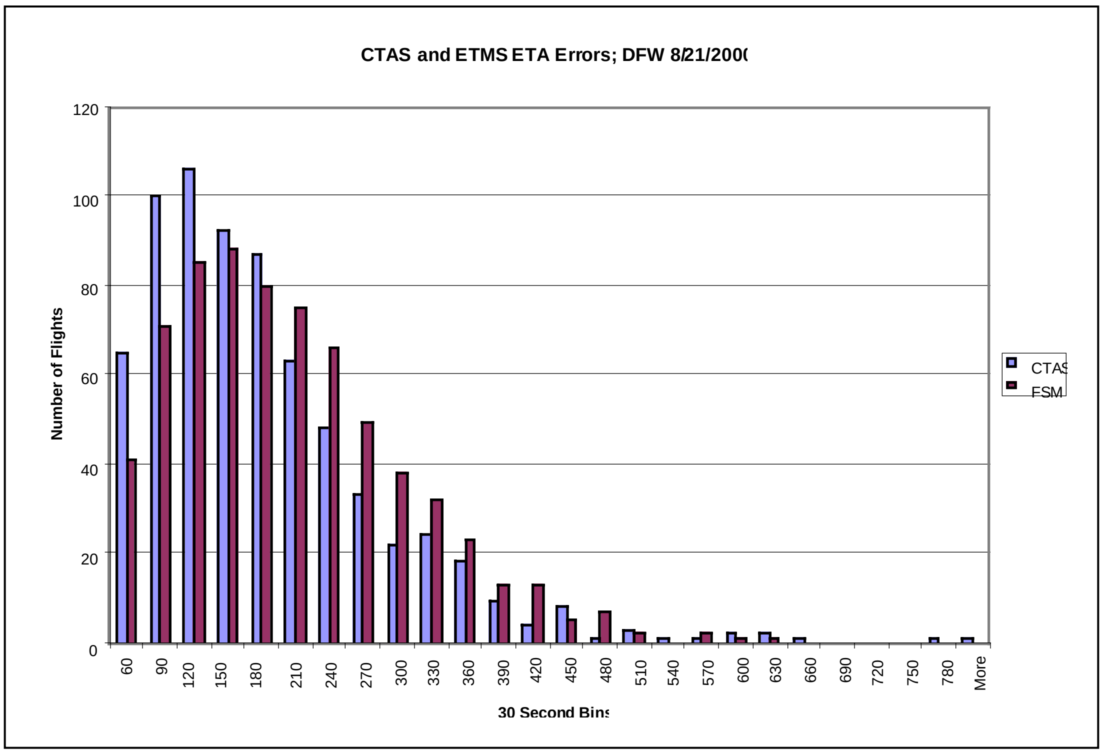
CTAS and ETMS ETA Errors; DFW 8/21/200(
120
100
80
60
40
20
0
Number of Flights
60
90
120
150
180
210
240
270
300
330
360
390
420
450
480
510
540
570
600
630
660
690
720
750
780
More
30 Second Bins
CTAS
FSM

Figure 4-2

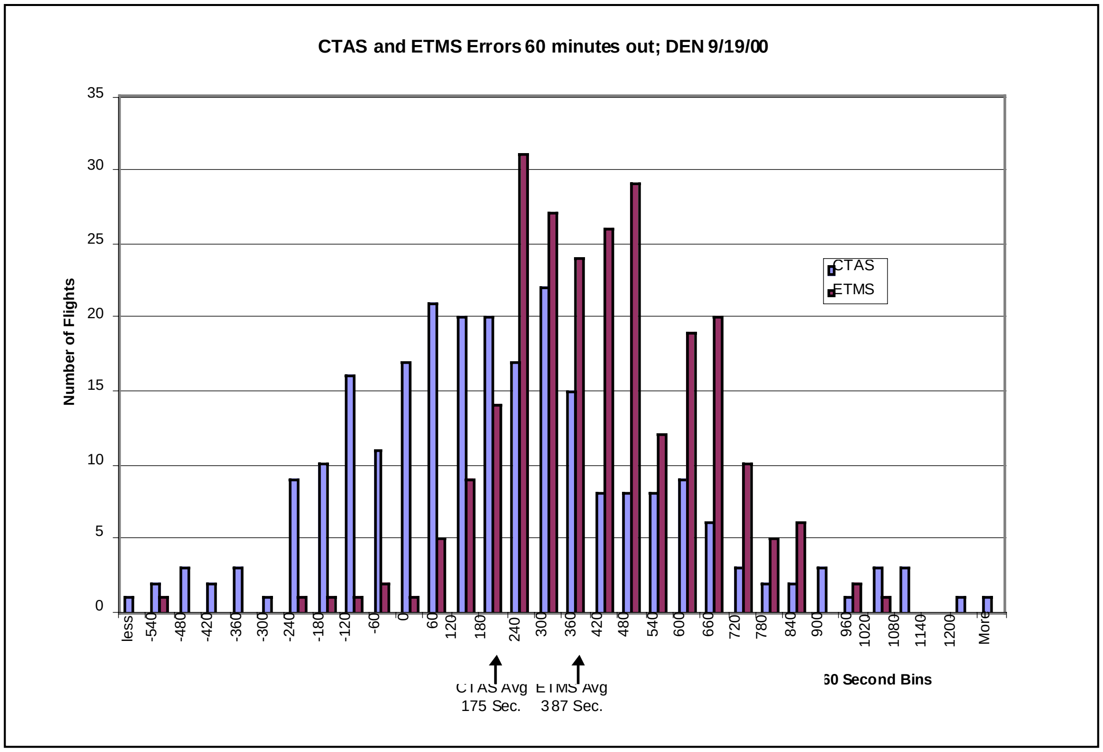

CTAS and ETMS Errors 60 minutes out; DEN 9/19/00
Number of Flights
35
30
25
20
15
10
5
0
CTAS
ETMS
less
-540
-480
-420
-360
-300
-240
-180
-120
-60
0
60
120
180
240
300
360
420
480
540
600
660
720
780
840
900
960
1020
1080
1140
1200
More
60 Second Bins
CTAS Avg
175 Sec.
ETMS Avg
387 Sec.

Figure 4-3

4.3 Hold/Go Decisions

Metric: Number of flight instances in which exactly one of the two systems (CTAS or ETMS) generated an ETA that fell within a specified range of the actual time of arrival of the flight and the other system fell outside the specified range.
Data: CTAS-simulated ETAs, ETMS ETAs.
AATT Category: Predictability

An inbound flight can delay a connecting flight if it is sufficiently late. At a certain critical time, the airline dispatcher must decide whether to hold the downstream flight for the delayed flight. This decision is known as a hold/go decision. A rule of thumb commonly used is to hold the downstream flight no more than 15 minutes beyond its scheduled departure time for the inbound flight. For instance, suppose that flight *f* is due to arrive at 10:15, connecting with flight *g*, due to takeoff at 11:00. Then the hold/go criterion is to hold *g* if *f* will arrive before 11:15 but to let *g* go if *f* will arrive after 11:15.

One of the benefits of improved arrival prediction accuracy is to allow an AOC to make effective hold/go decisions. This analysis is designed to capture that benefit. The metric we have established is to quantify those instances in which one system (CTAS or ETMS) gave an accurate arrival prediction for a hold/go decision but the other did not. For instance, ETMS could incorrectly predict that flight *f* would arrive 15 minutes or more after the departure time *T* of its connecting flight but CTAS could correctly predict that *f* would arrive within 15 minutes of *T*. There are four such cases in all, as shown in Table 4-4.

There are two complications associated with this metric. Without access to air carrier proprietary data, there is no way of knowing in historical hindsight whether or not a flight was connecting to an outbound flight. However, inbound and outbound banks of a hubbing operation can be strongly inferred by groupings in the schedule. Figure 4-4 shows the number of arrivals and departures at DFW between the hours of 1800z and 2200z on 10/25/00. The number of arrivals and departures for 15-minute periods are represented by juxtaposed bars; the lightly shaded regions represent the number of AAL arrivals, while the hashed, lightly shaded regions represent the number of AAL departures. The demand surges in arrivals and departures correspond to arrivals and departures of AAL banks. We found that the surges tend to be about 90 minutes apart. We used this type of informal approach to approximate bank periodicity.

The other complication is that we have no way of knowing how many minutes prior to arrival of an upstream flight an AOC would actually make a hold/go decision for a downstream flight. This might vary by airline and circumstance. Since we only have CTAS data for *T* values of 5, 10, .., 60 minutes prior to arrival, there is no point in selecting a value of *T* larger than 60. Choosing to err on the side of longer decisions rather than shorter, we selected *T* = 60 minutes.

Still, we cannot be sure that every flight we have selected is making a connection. Therefore, our metric will give an upper bound on the instances in which one system outperformed the other.

The metric was computed as follows. For each flight *f* in our data set, we assumed that it was connecting with a downstream flight *g* scheduled to depart 90 minutes after the scheduled arrival time of *f*. Scheduled arrival times were taken from the original estimated time of Arrival (OETA) field in the ETMS-based Delta files. Usually, this is the time listed in the Official Airline Guide but there are other sources, such as a daily schedule download from the air carrier to ETMS. The flight *f* is considered too late for connection if its actual arrival time was later than the departure time of *g* minus 15 minutes. The parameter 15 minutes was chosen as the minimum time that would be required at the hubbing airport for passengers to make their connecting flights. (The same might not be true for baggage and crews. Through informal experimentation, we later found that our results were relatively invariant with the value of this parameter.) Given our assumed parameter of 90 minutes between connecting flights, this means that *f* is considered too late if it arrives after the OETA of *f* plus 75 minutes (75 = 90 − 15).

Let

ETA_CTAS_f = the CTAS ETA made closest to 60 minutes before the arrival of *f*
ETA_CTAS_f = the ETMS ETA made closest to 60 minutes before the arrival of *f*

(If either ETA was missing, then *f* was excluded from the study.) Then ETA_CTAS_f is considered to form a *correct* basis for a hold/go decision if it falls on the same side of the $OETA_f + 75$ minute mark, that is if either of these holds:

(i) $arrival(f) > OETA_f + 75$ and $ETA\_CTAS_f > OETA_f + 75$
(ii) $arrival(f) \leq OETA_f + 75$ and $ETA\_CTAS_f \leq OETA_f + 75$

Otherwise, we say that ETA_CTAS_f is not correct. The correctness of ETA_ETMS_f is similarly defined. On a flight-by-flight basis, there are four possibilities:

- both ETAs are correct;
- neither ETA is correct;
- ETA_CTAS_f is correct and ETA_ETMS_f is not correct;
- ETA_ETMS_f is correct and ETA_CTAS_f is not correct.

The first and second cases indicate comparable performance of the two systems. The third case indicates better performance by CTAS. The fourth indicates better performance by ETMS.

The results of the metric are tallied in Table 4-5, Table 4-6 and Table 4-7. Unfortunately, there are very few instances in which either system prevails. Out of a total of 3,053 flight instances at DEN, there were

- 3,038 instances in which both systems were correct;
- 14 instances in which neither system was correct;
- 1 instance in which CTAS was correct (and ETMS was not);
- 0 instances in which ETMS was correct (and CTAS was not).

Thus, there are not enough prevailing flight instances (of the third or fourth kind) to draw any conclusions at DEN. Similarly, there were not enough prevailing instances at DFW or PHL to draw any conclusions, though ETMS prevailed in 7 instances and CTAS prevailed in none.

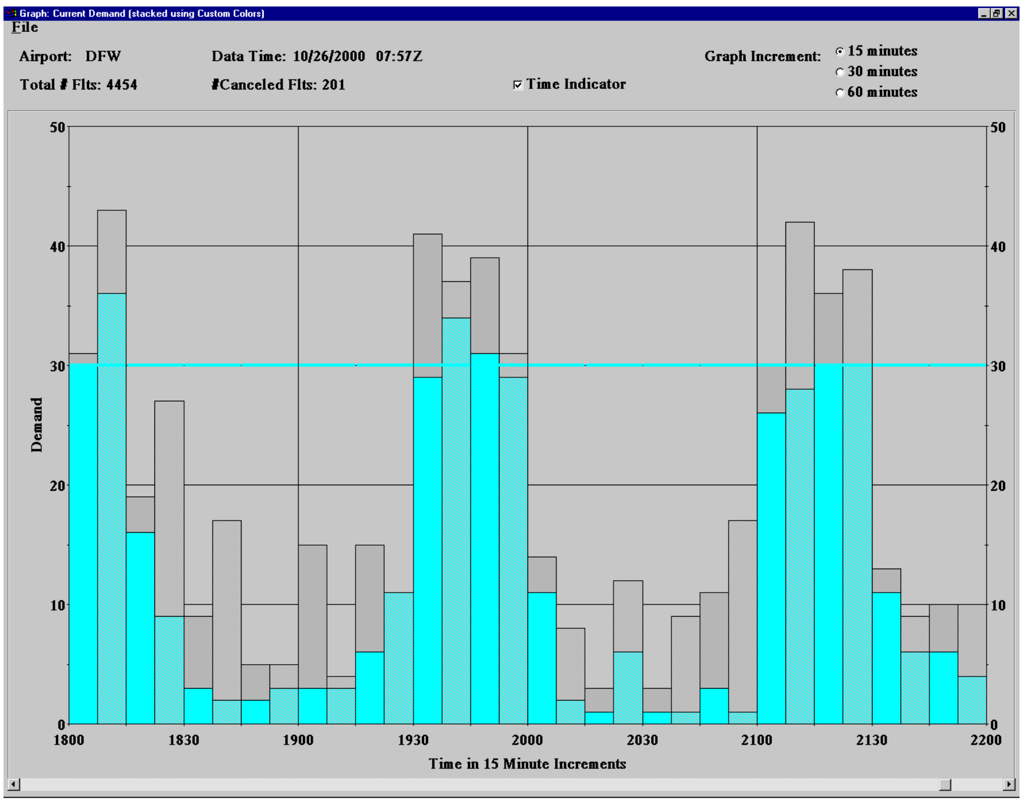

Figure 4-4

Table 4-4: Exhaustive List of Cases for Hold/go Decisions				
System	Predicted delay for outbound	Correctness	Decision	Case
CTAS	< 15	Yes	Hold	1
ETMS	≥ 15	No	Go	
CTAS	< 15	No	Hold	2
ETMS	≥ 15	Yes	Go	
CTAS	≥ 15	Yes	Go	3
ETMS	< 15	No	Hold	
CTAS	≥ 15	No	Go	4
ETMS	< 15	Yes	Hold	

DEN Hold/Go Decisions, CTAS vs. ETMS								
	9/14/2000	9/15/2000	9/16/2000	9/17/2000	9/18/2000	9/19/2000	9/20/2000	Total
Both	331	414	492	503	453	414	446	3053
Neither	3	1	5	1	3	0	1	14
CTAS	0	0	1	0	0	0	0	1
ETMS	0	0	0	0	0	0	0	0

Table 4-5

DFW Hold/Go Decisions, CTAS vs. ETMS								
	8/20/2000	8/21/2000	8/22/2000	8/23/2000	8/24/2000	8/25/2000	8/26/2000	Total
Both	770	730	795	761	783	740	751	5330
Neither	0	1	0	0	1	1	1	4
CTAS	0	0	0	0	0	0	0	0
ETMS	0	1	0	1	4	1	0	7

Table 4-6

PHL Hold/Go Decisions, CTAS vs. ETMS								
	9/5/2000	9/6/2000	9/7/2000	9/8/2000	9/9/2000	9/10/2000	9/11/2000	Total
Both	91	99						190
Neither	0	0						0
CTAS	0	0						0
ETMS	0	0						0

Table 4-7

4.4 Diversions

Metric: Number of flight instances in which exactly one of the two systems (CTAS or ETMS) generated an ETA that fell within a specified range of the actual time of arrival of the flight and the other system fell outside the specified range.

Data: CTAS-simulated ETAs, ETMS-generated ETAs.

AATT Category: Predictability

Flights are required by law to carry at least 45 minutes worth of extra fuel to provide for situations in which they encounter unexpected delay. As a rule of thumb, a flight will be diverted if holding goes beyond 45 minutes. In practice, the decision about whether or not to divert is highly case specific but we can make use of this rule of thumb to cite flights instances in which the need for a diversion was correctly forecasted.

We had planned to use a metric similar to that used in the hold/go decision analysis. The analysis would work as follows. For each flight *f*, set a time t = STA(f) + 45, where STA(f) is the carrier scheduled arrival time of flight *f*. In light of the fact that we have no diversion information available to us, this time t is a surrogate for the time at which a diversion decision must be made for flight *f*. Let $A(f)$ be the actual arrival time of *f*. Either $A(f)$ < STA(f) or $A(f)$ < STA(f). Flight instances would be tallied in which exactly one of the two systems correctly predicted this outcome.

This analysis is the same as the hold/go analysis except that we have added 45 minutes to the schedule arrival time of f instead of 75. Since the hold/go decision analysis was inconclusive, we did not pursue this approach for diversions.

4.5 Airborne Holding and Expected Clearance Times

Metric: IPE statistics for only those flights subject to airborne holding.
Data: Eastern Region Holding Statistics, CTAS-simulated ETAs, ETMS ETAs.
AATT Category: Predictability

When a flight is put into an airborne holding state (beyond vectoring or speed reduction) controllers give the pilot an expected clearance time, which is an estimate of the time at which the flight will be released from holding and sent to a metering gate at the airport [2]. Since this type of delay is dynamic and unpredictable, improved accuracy in clearance times is highly valuable.

We intended that the IPE statistics from Section 3.2 would be isolated for aircraft that were held. The concept was that for the flights that were held, the CTAS predictions of ETA taken just prior to holding will be significantly better than those of ETMS.

Though the concept here is sound, the availability of reliable holding statistics a prevented us from applying it. The dates on which CTAS data was collected did not overlap with the dates for which we had holding data. Also, the holding data is currently limited to the eastern region of the country and would have covered only two days at PHL. Moreover, the CTAS collection dates were fair weather days and would not likely be days on which significant amounts of holding took place.

5 Benefits Assessments Related to User Preferences

5.1 User Preference Input to CTAS Scheduling via CAP

In this section, we assume the existence of CAP function to receive inputs from air carriers as to the relative delay costs associated with their flights. These costs would act as a guide for arrival time adjustments and performing tradeoffs (delay exchanges) between flights of the same carrier or of differing carriers.

We have assumed the input of flight cost functions that represent the priorities and preferences of the carriers. Once we discretize time into t-many slots, we can compute the cost of assigning each flight f to each of the slots $t = 0, 1, \ldots, T$. These values can be computed in advance and stored in a table for reference.

We envision that the flight cost structures would be relatively simple, probably in the form of discrete delay levels, as in Figure 5-1. The motivation for cost structures is the recognition of operationally significant delay levels within each carrier. For instance, the industry standard for an on-time arrival is 0 to 15 minutes of delay beyond carrier scheduled arrival time. Thus, the difference between 7 and 9 minutes of delay is not nearly as significant to the carrier as the difference between 14 and 16 minutes of delay. A common carrier priority is to decrease the number of flights that have more than 15 minutes of delay. Similarly, when delay levels are between 15 and 25 minutes, the rate of missed baggage connections begins to increase. Between 25 and 45 minutes of delay, passengers begin to miss connections. And over 45 minutes of delay, crews begin to have a significant rate of missed connections, along with passengers and baggage. This develops the significant delay levels, as in Figure 5-2. Each carrier would have to set the upper and lower limits of these delay 'buckets' based on the characteristics of their operation and network connectivity. The delay buckets could be flight specific or set uniformly for all flights within an operation.

The weights (delay costs) associated with each delay level would be arbitrary units used for relative comparisons of delay costs. Flights with higher priority would be given higher weights. Normalization must be performed, of course, to prevent a carrier from gaming the process by entering large cost values. The exact manner in which normalization is conducted depends on the manner in which equity is addressed. One method for normalizing is to divide all costs for a given carrier by the maximum cost supplied by that carrier. This has the effect of placing all costs for the carrier on a scale of 0 to 1.0.

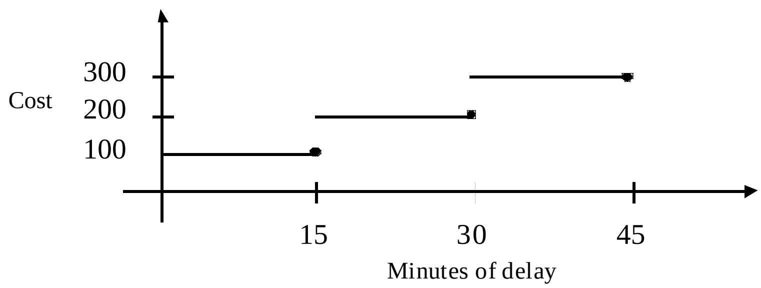

Figure 5-1: A hypothetical flight cost function

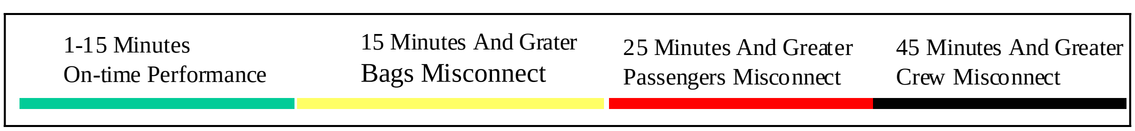

Figure 5-2: Potential delay buckets

5.2 The Concept of Delay Exchange

At the time of this study, it is uncertain exactly how the air carriers will input preferences to CTAS scheduling via the CAP tool. However, it is reasonable to assume that they will be most intent on maintaining the integrity of their schedule. Once a flight has entered the CTAS domain of operation, it may have already assumed delays due to a combination of late departure or enroute restrictions. Traffic congestion and physical limitations on speed increases severely limit the ability of air traffic service providers to offset this delay post facto.

Also, in the aggregate, any delay of aircraft suggested by CTAS to air traffic service providers for smoothing out local demand surges is also unavoidable. But prior studies have shown that much can be done toward prioritizing flights once delays are within the CTAS realm [7] [8]. The basic idea is to delay a flight running early or on schedule in favor of a more critical flight running late.

There are two ways that CAP may incorporate tradeoffs between flights.

- *Intra-carrier delay trading*, in which each air carrier would tradeoff, anticipated delays within its own set of flights so as not to interfere with flights of other airlines. In practice, such a process would involve a combination of automation and treatment of flights on an individual basis at the airline operational centers.
- *Inter-carrier delay trading*. The tricky part of this is to make realistic assumptions about how equity between carriers would be achieved. This requires balancing what may be conflicting requests from carriers for speed reductions or increases.

Under contract with NASA, Metron developed an Equity Swapping Algorithm (ESA). ESA is an automated procedure for performing arrival time adjustments in which flights (possibly between differing carriers) effectively trade the original arrival times assigned to them by CTAS. For this reason, it can be thought of as an inter-carrier slot swapping algorithm, comparable to the intra-carrier slot swapping mechanisms used by carriers to monitor their own operations during one of the FAA's ground delay programs. In the case of ESA, delay intervals, or *buckets*, were used as the criteria for slot swapping, as described in Section 5.1.

We sought to identify the instances in which flights could be sped up enough to qualify for a lower delay bucket while leaving all other flights in their original delay buckets. Essentially, this will be done by swapping flights between STA arrival slots created by CTAS.

We modeled the CTAS boundary of operation by a fixed time radius of T = 60 travel minutes out from the airport. For each flight in our study, we used ETMS-generated historical data to capture the ETA of the flight (ETA_{60}) as it crossed the hypothetical CTAS boundary. We defined the delay of a flight as it crosses into the CTAS realm as $ETA_{60} - OETA$, where $OETA$ is the original estimated time of arrival of the flight, as recorded by CDM/ETMS. In most cases, this is the arrival time scheduled by the air carrier, as published in the Official Airline Guide (OAG).

For each of three airports, ORD, DFW and PHL, we selected one or more air carriers with a dominant presence at the airport and performed a rescheduling of arrival times on their behalf, in order to hypothetically model a flight prioritization scheme. We assumed that the airline had the ability to input to CTAS preferences for arrival times subject to the following constraints:

1. flights are reviewed for delay by the airline and processed by CTAS sequentially, as they cross over the CTAS boundary;
2. arrival time adjustments must lie in the range [−5,10] minutes;
3. the total delay incurred by each airline must remain constant;
4. the adjusted arrival time of each flight must correspond to STA_{60} for some flight belonging to the same airline.

The first constraint reflects the fact that an entire arrival order cannot be considered at one time; arrival time adjustments must be made dynamically, as information becomes available. This is in direct contrast to the strategic arrival slot allocation and manipulation that is performed in planning a ground delay program (GDP).

The second constraint is consistent with the bounds on arrival time adjustments we used in our prior analyses and are our best approximation to a uniform adjustment window for all flights.

The third and fourth constraints preserve equity among competing airlines. In essence, the airline is considered to have control over (own) a set of arrival slots based on STAs captured upon entry to the CTAS realm. This is analogous to the ration-by-schedule algorithm established by the CDM community that initially allocates arrival slots to airlines at the start of a GDP, based on OETAs (see [3] or [6] for a reference on CDM and GDP-E procedures). The fourth constraint says that each flight must be substituted into an arrival slot owned by the same airline.

Moreover, we assume that all requests for arrival time adjustments that are feasible with respect to the above criteria will be honored by air traffic service providers. Since not all could be honored, our results are an upper bound on the potential for arrival prioritization by delay increments.

Based on the guidance of an experienced (former) airline dispatcher, we established the set of delay buckets in Table 5-1.

Table 5-1: Delay Bucket Parameters			
Delay Bucket	Lower Bound (min)	Upper Bound (min)	Operational Sensitivity Issue
B_1	0	14	On-time performance
B_2	15	20	Baggage connection
B_3	21	40	Passenger connection
B_4	41	50	Crew connection/legality
B_5	51	60	Equipment shortage
B_6	61	90	Aircraft maintenance
B_7	91	Inf	Flight cancellation

Flights were processed in order of CTAS boundary arrival (inferred by increasing order of ETA_{60}), reflecting the fact that they would most likely be considered for arrival time adjustment on a first-come, first-served (FCFS) basis. Our analysis breaks away from current FCFS practices in that it allows for instances in which the relative ordering of flights entering the CTAS realm may differ from the ordering of arrivals on the runways.

We adopted a simple strategy for processing airline preferences: whenever a flight *f* is in a delay bucket other than B_1, the on-time bucket, an attempt was made to swap the arrival time of *f* with the arrival of another flight *g* ahead of it in the CTAS-boundary arrival queue. This preserves total delay because it adds to *g* the delay saved by *f*. However, the swap is made only if

(1) flight *f* can be moved to a lower delay bucket and
(2) flight *g* remains in its current delay bucket.

Then net the effect of the swap is a movement of one flight from delay bucket B_n to delay bucket B_{n-1}, where B_n is the original delay bucket of *f*.

A typical swapping scenario would be as follows. Consider the ETAs listed as in the table below.

Flight	ETA	STA	Earliest Slot
A_1	1210	1215	1205
B_1	1215	1220	1210
A_2	1220	1225	1215

Suppose that airline A wishes to move one of its flights, A_2, into an earlier slot. Given its ETA of 1220 and a maximum speed up of 5 minutes, the earliest slot A_2 can assume is 1215. Unfortunately, the 1215 (ETA-based) slot is owned by competitor *B*, so no advancement is possible. The TMA component of CTAS routinely suggests to air traffic service providers the addition of delay to select flights whenever demand exceeds capacity of the airport. At some airports, airborne holding routinely hovers around 5 to 15 minutes per flight at busy times of the day. This effectively adds to the number of minutes that a flight can be sped up. Let us return to our hypothetical example and suppose that CTAS is recommending delays of five minutes per flight. Prior to airline intervention, the flights will have scheduled times of arrival listed in the STA column of the table. Since it is still feasible for A_2 to arrive at 1215, A_2 can trade (STA-based) arrival slots with flight A_1; a swap is achieved where none was possible before. In effect, scheduled airborne delays add to the number of minutes upward (earlier) in the schedule that a flight can reach, when considering a swap. (Similarly, in a GDP, slot swapping is enabled by FAA-assigned delays.) This operational behavior of CTAS is the essential factor that enables slot swapping on a regular basis.

Arrival time adjustment is limited by how much a flight can be physically sped up or slowed down in the latter stages of flight. Early experiments performed at Metron revealed that there is little potential for swapping, based on the ability to adjust the arrival time of a flight by only five minutes [8]. However, there is considerable potential for delay swapping once simulated demand smoothing is introduced. Just before each flight was processed for airline prioritization, an initial adjustment was made to its ETA to yield a CTAS scheduled time of arrival (STA) based on both airport capacity and the STAs of the flights ahead of it. Slot swapping will then be performed based on these temporary STAs to arrive at final STAs.

We used the following airline-airport pairs.

Code Name	Airline	Airport
UAL-ORD	United Airlines (UAL)	Chicago O'Hare (ORD)
AAL-ORD	American Airlines (AAL)	Chicago O'Hare (ORD)
AAL-DFW	American Airlines (AAL)	Dallas-Fort Worth (DFW)
USA-PHL	US Airways (USA)	Philadelphia (PHL)

UAL-ORD and AAL-DFW have been chosen to represent large hubbing operations; USA-PHL was chosen to represent a smaller hubbing operation; AAL-ORD was chosen to represent a large airline at an airport that is not its primary hub. For each operation, and for each of five non-GDP days, we applied the combined CTAS and airline prioritization simulator to all flights that arrived at the airport. Days on which a GDP had been implemented were avoided because, during a GDP, most flights are arriving in accordance with FAA-issued arrival times, rather than in accordance with their natural variations from scheduled arrival times. Essentially, we used the non-GDP status as a surrogate for fair weather conditions and absence of widespread, FAA-imposed traffic restrictions.

5.3 Inter-carrier Delay Exchange via Slot Swapping

Metric: Increased on-time arrivals, delay minutes/dollars saved
Data: CTAS-simulated ETAs, ETMS ETAs.
AATT Category: Flexibility

For select days at select airports, we ran ETMS flight data through our CTAS scheduling and delay exchange simulator to assess the potential for delay exchanges that could occur if managed by the CAP tool. Our approach was to compute upper and lower bounds on this potential. In the lower bound computations, we placed the following two restrictions.

- R1: *No passing flights within the same stream class.* An arrival stream class is determined by similarity of aircraft type (e.g., jet or turbo-prop) and arrival fix used. In addition, when a flight was moved into an earlier arrival slot in a delay exchange, we required that there be enough spacing between a flight and the flights in front of it in its arrival stream to perform the speed increase. This restriction prevents a situation in which an aircraft would be physically required to pass other flights within its arrival stream.
- R2: *Equity in delay exchange*: Essentially, this restriction says that at the end of each day, the number of minutes given up by a carrier to its competitors must be equal the number of minutes assumed from its competitors. We used number of swaps as a surrogate for delay minutes exchanged. That is, if carrier A had 10 swaps in which its flights assumed more delay, then it must have 10 swaps in which it flights had reduced delay. See Section 62, "Concept of Delay Exchange" for further details.

The assumption here is that for an airport with four arrival fixes, for instance, the traffic over each arrival fix would be divided by aircraft type for a total of eight distinguishable stream classes. We comment that we cannot be guaranteed that for every flight the arrival fix we obtained in the ETMS database was the one the flight actually used. In some instances, this could be the *filed* arrival fix, as opposed to the actual. However, it seems reasonable that these deviations cancel each other in the aggregate so that we still had a representative cross-section of flights of each aircraft type over each arrival fix.

Upper bound results for a given day at a given airport were computed by simulating delay exchanges without restrictions R1 and R2. Lower bound results were computed by simulating delay exchanges with restrictions R1 and R2. First, we discuss the upper bound results.

Table 5-7 summarizes the delay exchanges we simulated using our ORD 6/18/99 ETMS data set. For now, we restrict our attention to the row marked "All", which indicates that inter-carrier delay exchanges were allowed to take place between any two carriers. Out of 1224 flights arriving at ORD, we found that there was a potential for 149 inter-carrier delay exchanges (swaps). The Max Swap Column shows that some of these delay exchanges required arrival time adjustments as high as 15 minutes. This does not mean that aircraft are sped up or slowed down 15 minutes. In our simulation, delay is added to flights to simulate demand smoothing that would be recommended by CTAS. This delay then becomes “negotiable” in the exchange process.

The maximum number of position shifts caused by the reordering of the arrival hierarchy is given by the Max Flights Passed column: 21 flights. We have reason to believe this is feasible because the flight does not physically pass other flights, it merely assumes a different position in the list of flights when ordered by CTAS-schedule arrival times. The average number of minutes of delay exchange was 3.65 minutes.

The most remarkable feature of Table 5-7, and perhaps the most remarkable potential benefit of CAP, lies in the columns marked $B_2 \rightarrow B_1$, $B_3 \rightarrow B_2$, etc. Each of these columns indicates the number of flights that were moved from one bucket to another. For instance, the value 71 in row “All” and column “$B_2 \rightarrow B_1$” tells us that in the inter-carrier delay exchange process (All), there were 71 flights moved from delay bucket B_2 into the on-time category, bucket B_1. That’s 71 more on-time arrivals than would have taken place under the first-come first-served (FCFS) scheduling practices currently used. Similarly, the number of movements between the other delay buckets were as in Table 5-2.

Table 5-2							
Bucket Swaps at ORD 6/18/99							
$B_2 \rightarrow B_1$	$B_3 \rightarrow B_2$	$B_4 \rightarrow B_3$	$B_5 \rightarrow B_4$	$B_6 \rightarrow B_5$	$B_7 \rightarrow B_6$	$B_3 \rightarrow B_1$	$B_4 \rightarrow B_2$
71	48	12	10	6	1	1	0

This demonstrates the enormous potential that CAP has for mitigating damages incurred by flights falling behind schedule.

These results assume maximum participation by the air carriers in a delay exchange program, as well as complete ability of air traffic service providers to accommodate all profitable delay exchanges. For this reason, the results above are truly an upper bound on savings due to delay bucket reductions.

Table 5-8 shows results comparable to those in Table 5-7 only for a different day at ORD (4/19/99). This just exemplifies our informal observations that results were fairly consistent from day to day.

5.3.1 Summary Results for Inter-carrier Delay Exchanges

We ran delay exchange (slot swapping) simulations at ORD, DFW, and PHL for each day of the first six months of 1999. Some days were excluded from the study, for one of two reasons: either a ground delay program (GDP) was conducted by the FAA, or the arrival count was so abnormally low that we concluded that the data file was incomplete.

Table 5-3 gives the daily average results of the Inter-carrier slot swapping simulation. Averages were computed for each day, and then averaged over all days. (An alternative approach is to weight the daily averages by the number of flights each day before averaging.) Column A gives the number of swaps that occurred each day. A single flight may have been involved in more than one swap but an informal examination revealed that this was a rare occurrence. Column B gives the average number of minutes that was traded in a swap. This is the number of minutes by which one flight would be delayed and the other sped up. Column C gives the number of flights that were moved into the on-time arrival bucket (B_1) as a result of the swaps. The values in this column are rather impressive, considering that an increase of only a few on-time arrivals per day can make a significant difference in carrier performance and passenger satisfaction. The values ranged from 29.22 more on-time arrivals per day at PHL to 57.85 at DFW.

The percentage *increase* of on-time arrivals was computed by taking the difference of the percentage of on-time arrivals before swapping (on-time arrivals divided by total flights) and dividing it by the percentage of on-time arrivals after swapping. The lower bound is 1.81% increase, while the upper bound is 5.77% increase. (See Table 6-1).

We comment that there would be substantial additional savings to an inter-carrier slot swapping program. In addition to flights being moved from delay bucket B_2 to the on-time bucket B_1, many flights were also moved from B_3 to B_2, from B_4 to B_3, and so on. We were not able to associate any dollar amounts with these savings. This is in part because it would require proprietary airline information, which we are obligated to avoid in this study. Also, we tried to get educated estimates from airline personnel on these values but, in many cases, the cost of a higher delay bucket came out less than the cost of a lower delay bucket. This would mean that rather than move flights to lower delay buckets whenever possible, as was done in our simulation, we should have intentionally delayed some flights to move them up into a less expensive delay bucket. This struck us as counter-intuitive, or even contradictory. Since this type of cost structure was inconsistent with the approach used in our simulation, we abandoned any further attempt to quantify the cost savings associated with higher delay buckets. Moreover, these costs would be highly airline-specific and situation-specific.

In order to compute a lower bound on the potential value of an inter-carrier slot swapping program, we repeated the simulations already described over the same sets of flights but with the added restrictions R1 and R2. Lower bound results are listed in the lower rows of Table 5-3.

Note that there is a significant gap between the upper and lower bounds. Since the number of minutes exchanged in a slot swap tends to be close to the average, the total number of swaps that occurs serves as a rough guide for the benefits received. At ORD, the upper and lower bounds on average swaps per day were 125.01 and 23.21 respectively. We were concerned that restrictions R1 and R2 would reduce the number of swaps to very few per day. However, figures close to the lower bound, 23.21, still produce remarkable daily and annual savings.

In the summary table at the end of this document (Table 6-1), we have listed the lower bound results as our best estimate. There are two reasons for this. One is that it seems reasonable that the "no-passing" and "equity" restrictions R1 and R2 would need to be enforced. The other reason is that one of our blanket assumptions in this study is that air traffic service providers could, and would, carry out an advisories posted by the CAP tool. In practice, this may not be the case. One possible area of future study would be to examine the workload that would be placed on air traffic service providers by a delay exchange program, and how often the delay exchanges could actually be carried out.

Since the number of swaps is roughly proportional to the number of arrivals (traffic demand), this percentage can be extended to NAS-wide implementation without referring to numbers of operations across the NAS now, or in the future. This is how we arrived at the corresponding increases in on-time performance in our final summary Table 6-1 (cells F5 and G5). Below each percent increase are a number of flights. This was computed NAS-wide by multiplying the percentage increase by the number of arrivals in 1999 for the top 38 airports (see Table 8-1). That is,

Lower bound = 1.81% of 8,449,220 flights = 152,931 flights
Upper bound = 5.77% of 8,449,220 flights = 487,520 flights

These two underlined figures have been entered in Table 6-1 in cells F6 and G6, respectively.

Multiplying by an average of 4.08 minutes per swap and by $49.90 savings per minute, we obtain the following bounds on annual savings.

Upper bound = 487,520 flights × 3.96 min per swap × $49.90 per min = $95,362,812
Lower bound = 152,931 flights × 3.96 min per swap × $49.90 per min = $29,914,527

These two underlined figures have been entered in Table 6-1 in cells F7 and G7, respectively. Annual savings in dollars were computed as in Section 6.3 and entered into cells F10 and G10 of Table 6-1.

Table 5-3: Inter-carrier Slot Swapping Daily Averages

	A	B	C	D
Airpt	Number of swaps	Swap length in minutes	Increased number of on-time arrivals	Increase in on-time performance (%)
	Upper Bounds			
ORD	125.01	3.79	57.85	5.37
PHL	64.07	4.46	29.22	5.42
DFW	127.41	4.00	39.62	6.53
Avg.	105.50	4.08	42.23	5.77
	Lower Bounds			
ORD	23.21	3.22	35.36	3.29

PHL	10.58	4.09	5.34	0.99
DFW	60.42	3.86	11.32	1.14
Avg.	31.40	3.72	17.34	1.81

5.4 *Intra-carrier Delay Exchange via Slot Swapping*

Metric: Increased on-time arrivals, delay minutes/dollars saved
Data: CTAS-simulated ETAs, ETMS ETAs.
AATT Category: Flexibility

Many of the delay exchanges in the inter-carrier delay exchange of Section 5.3 are actually intra-carrier exchanges, e.g., a UAL flight swapping arrival slots with another UAL flight. Complete participation by all the carriers is dubious; not all carriers would have the capability or desire to participate. For this reason, we assessed the potential for intra-carrier slot swapping for select carrier list in the first row of the table in Figure 5-3. Note that intra-carrier swaps were possible only for the two major carries at ORD, AAL and UAL. On this particular day, these carriers had 543 and 463 arrivals, respectively. This represents approximately 82% of the total number of arrivals.

As one would expect, the number of feasible swaps is down from 149 (when all carriers participate) to 47 for UAL and 41 for AAL. That is, the number of swaps is reduced to slightly less than one-third of the or its original amount. For each of these carriers, the distribution of bucket movements is similar to the distribution of bucket movements when all carriers participate except that it is greatly reduced. Still, this shows enormous potential for delay bucket reductions for the major or hubbing carriers. For instance, UAL could have 23 more on-time arrivals at ORD on this day, while AAL could have 26.

5.4.1 Summary Results for Intra-carrier Delay Exchanges

The simulations and methods of computation for benefits that would result from a single carrier participating in a delay exchange program were identical to those of an inter-carrier delay exchange program, except that swaps were limited to a single carrier. We primarily relied on UAL and AAL operations at ORD for our single carrier scenarios. We caution that these are hubbing operations. Benefits for non-hubbing carriers can be close to zero, as evidenced in Figure 5-3. We feel that this is a reasonable representation because it is understood that in the single carrier setting, only a large hubbing airline would have enough flights present to make effective swaps. (If flights are spaced too far apart, there is no opportunity for slot swapping.)

We chose several representative days at ORD: 6/18/99, 4/19/99, and 5/28/99 and averaged the percentage increases in on-time arrivals (flights moved into bucket B1). See Table 5-4. We found the average increase in on-time performance to be 5.05%. Note that this is slightly lower than the value arrived at for the inter-carrier delay exchanges (5.77%). This makes intuitive sense, since the large hubbing operations were receiving most of the benefits anyway. Since the figures in Table 5-4 were from a simulation in which restrictions R1 and R2 were not in effect, they are upper bounds on benefits. Comparable computations were made for lower bounds to arrive at a 2.01% increase in on-time performance. The figures 5.05% and 2.01% were entered into the final summary Table 6-1, in cells G8 and F8, respectively.

The increase in the number of on-time arrivals (cells F9 and G9) in Table 6-1 were computed by taking 2.01% and 5.05% of the average number of flights for a hubbing operation (Avg. flights from Table 5-4) and multiplying by 285 non-GDP days per year. The respective figures were 2,855 and 7,172 flights per year. We chose the lower bound as the best estimate, for reasons already stated in Section 5.3.

The annual savings of

Upper bound = $558, 399
Lower bound = $1,402,943

were computed in a manner similar to those in Section 5.3 and entered into cells F10 and G10 of Table 6-1.

Table 5-4: Increased On-times for Single Carriers				
Upper Bounds				
Date	Airline	Flights	Increased No. On-Times	Percentage Increase in On-time Performance
6/18/1999	UAL	543	23	4.24
	AAL	463	26	5.62
4/19/1999	UAL	560	30	5.36
	AAL	454	25	5.51
5/29/1999	UAL	529	25	4.73
	AAL	441	22	4.99
	Avg.	498.33	25.17	5.05

5.5 Arrival Time Adjustments by Air Carrier Preferences for Schedule Preservation

Metric: Increased on-time arrivals, delay minutes/dollars saved
Data: CTAS-simulated ETAs, ETMS ETAs.
AATT Category: Flexibility

A prime concern of many of the air carriers is the preservation of schedule integrity, particularly when arrivals are being constrained due to limited capacity or the type of local smoothing of demand surges that would be done by CTAS. We anticipate that the CAP tool will be able to accept preferences from the AOCs that are designed to restore their (air carrier) schedules as much as possible. For an individual flight, schedule integrity means setting the STA (CTAS-scheduled time of arrival) as close to the air carrier schedule time of arrival as possible. Assuming maximum scheduling efficiency on the part of CTAS, this may mean having to adjust the arrival times of other flights to accommodate. A more conservative approach is to allow only those arrival time adjustments that would not jeopardize spacing requirements with other flights within the same arrival fix.

In order to assess the potential for pure delay reductions, we ran the delay exchange simulation of Section 5.3 but allowed flights to move only into a vacant arrival slot rather than into a slot occupied by another flight. This can be thought of as a delay exchange with a phantom flight, or as a pure arrival time adjustment. We have assumed, of course, that the TMA component of CTAS has been supplied with up-to-date schedule information, and that air traffic service providers are successfully able to work towards a schedule output by CAP.

The bottom rows of the tables in Table 5-7 and Table 5-8 give the results of this analysis for 6/18/99 and 4/19/99, respectively. On 6/18/99 at ORD, we found that over all the carriers, there was a potential for 60 instances in which a flight could be sped up to move it into a lower delay bucket. In particular, there were 29 instances in which a flight could be moved into the on-time category. This could have a substantial impact on carrier performance statistics. Comparable statistics were found at ORD on 4/19/99; see Table 5-8.

5.5.1 Summary Results

As in Sections 5.3 and 5.4, we ran our simulation on each day's data from ORD, DFW and PHL for the first six months of 1999, excluding GDP days and days when data was corrupted. The summary results for the six-month period are in Table 5-5.

In Table 5-5, the product of Column A and Column B is the number of minutes saved by arrival time adjustments. To obtain the average daily savings, we multiplied this figure by the operating cost $49.90 per minute. Thus, the average daily savings, Column E, was computed as follows:

$$\text{Column E} = \text{Column A} \times \text{Column B} \times \$49.90/\text{minute}$$

The average airborne operating cost of $49.90 per minute ($2,994 per hour) was taken from Column 8 of Table 4-10 in the document "Economic Values for Evaluation of Federal Aviation Administration Investment and Regulatory Programs". This document was prepared by the Federal Aviation Administration's Office of Aviation Policy and Plans (APO) [11]. This estimate is based on weighted operating costs for scheduled, commercial service (Table 4-10 of the document has operating costs for other classes of flights as well).

An alternative figure that could be used for airborne operating cost is the Air Transport Association's estimate of $48.10 per hour. The ATA and APO figures are fairly close; the technical monitor at NASA requested that we use the figure from APO. Where desired, the reader can substitute the ATA value to arrive at a final dollar amount.

5.5.2 NAS-wide Results

In order to estimate the savings that would results from NAS-wide implementation of the CAP tool, we extrapolated the savings at DFW, ORD and PHL to the 38 airports in the NAS with the most number of operations. Table 8-1 is a listing of the airports that were in the Top 38 of all airports in the NAS for number of operations in the year 1999 or the year 2015. Column A is the number of operations for year 1999, while Column B is the (projected) number of operations for year 2015. Since two of the airports were in the top 38 for one year but not the other, there are 40 airports listed, rather than 38. Row 41 is the sum of the operations over just the top 38 airports for the respective years. Thus, this is slightly less than the sum of the figures above. An operation is either an arrival or a departure. Row 42 is the total number of operations divided by 2. We used the figures from (Row42, ColA) and (Row42, ColB) to project our results over the Top 38 airports. Although this is not the entire NAS, we will refer to this as "NAS-wide" results. It is dubious that CTAS, let alone CAP, would be deployed across the NAS. The main intent of this approach is to put the benefits of the CAP tool in context with other AATT studies that have been done. These studies have computed benefits for deployment at the top 38 airports.

As stated earlier, there is little deviation in the average number of minutes per swap across our three test-case airports ORD, DFW, and PHL. The average is close to 4.0 minutes. The difficult part of this extrapolation is estimating the number of swaps that would take place at other airports. For this, we needed to write a formula for swaps as a function of arrivals. We proceeded as follows for the lower bound case. At PHL, there was an average daily arrival count of 541.56 flights. On average, there were 10.58 swaps per day, where by "swaps" we mean swaps into vacant arrival slots. By dividing the latter figure into the former, we get a flight-to-swap ratio of

(PHL) 541.56 / 10.68 = 50.71 flights per swap

This can be interpreted to mean that there is one swap for every 50.71 flights. Similar computations were made for ORD and DFW:

(ORD) 1075.95 / 23.89 = 45.04 flights per swap
(DFW) 991.99 / 12.05 = 82.32 flights per swap

We averaged these three values to arrive at a single flights-to-swap FSR ratio:

FSR = 59.36

There were 8,449,220 arrivals in the NAS in 1999. Dividing by 365 days per year, this implies 23,148.55 arrivals per day. Dividing by our flights-to-swap-ratio (FSR) of 59.36, this implies 389.97 swaps (flight adjustments) per day. Multiplying by our average of 3.92 minutes per swap, we obtain 1,528.68 minutes per year. Multiplying by the FAA-APO cost of $49.90 per airborne minute, we obtain $76,281.02 dollars per day. The three underlined figures we have just computed were entered into Table 5-5 in the row marked "NAS". To project this final daily dollar figure over the entire 1999 year, we multiplied by a conservative estimate of 285 non-GDP days per year to obtain a total savings in 1999 of $21,740,090. These last three underlined figures have been entered into Table 5-6 in the row marked "NAS".

This computation was repeated for year 2015 except that instead of beginning with the 1999 number of operations (8,449,220), we used the year 2015 number of operations (12,048,417), taken from Table 8-1. Thus, the projected dollar savings due to arrival time adjustments is projected to be $31,000,929. This figure is, of course, in 1999 dollars.

Since there is no notion of equity in this analysis, the annual dollar figures we have just quoted, $21,740,090 and $31,000,929, have been entered into the final summary Table 6-1 as the upper bound, the lower bound, and the best estimate.

The percentage increase in on-time arrivals and the increased number of on-time arrivals in Table 6-1 were computed as follows. The increase in on-time percentage was extrapolated from the average in Column D of Table 5-5. Since this is a percentage, it requires no adjustment for operations per year. The increased number of on-time arrivals is simply 2.06% of the total operations for year 1999 (8,449,220). For year 1999, this is 174,054. For year 2015, this is 248.197.

Table 5-5: Arrival Time Adjustment Daily Averages					
	A	B	C	D	E
Arpt	Number of adjustments	Adjustment in minutes	Increased number of on-time arrivals	Increase in on-time performance (%)	Savings (at $49.90 per minute)
ORD	54.40	3.90	23.89	2.28	$10,586
PHL	27.08	4.23	10.68	2.05	$5,716
DFW	48.00	3.63	12.05	1.85	$8,695
Avg.	43.16	3.92	15.54	2.06	$8,332
NAS	390	3.92	174,054	2.06	$76,281

Table 5-6: Arrival Time Adjustment, Savings for Year 1999			
	A	B	C
Airport	Daily Savings per non-GDP day (E, Table 5-5)	Non-GDP Days in 1999 (B, Table 8-2)	Savings for 1999 A × E
ORD	$10,586	285	$3,017,010
PHL	$5,716	322	$1,840,552
DFW	$8,695	362	$3,147,590
Avg.	$8,332	323	$2,668,384
NAS	$76,281	285	$21,740,090

ORD 06-18-99; AAR = 90															
	Number of Flights	Number of Swaps	Avg Swap Mins	Max Swap Mins	Max Flights Passed	Max Swaps in Hour	Bucket Swaps B2-B1	B3-B2	B4-B3	B5-B4	B6-B5	B7-B6	B3-B1	B4-B2	Diff AC Type & F
ALL	1224	149	-3.65	15	21	17	71	48	12	10	6	1	1	0	111
UAL Only	543	47	-4.13	13	9	8	23	16	1	4	1	0	2	0	33
AAL Only	463	41	-3.93	12	7	7	26	7	2	4	1	0	1	0	30
COA Only	21	0	0.00	0	0	0	0	0	0	0	0	0	0	0	0
DAL Only	28	0	0.00	0	0	0	0	0	0	0	0	0	0	0	0
NWA Only	28	0	0.00	0	0	0	0	0	0	0	0	0	0	0	0
USA Only	19	0	0.00	0	0	0	0	0	0	0	0	0	0	0	0
SLOT Swaps	1224	60	-3.53	-1	11	6	29	19	8	2	2	0	0	0	

Table 5-7: Bucket Swaps at ORD, 6/18/99

ORD 04-19-99; AAR = 90															
	Number of Flights	Number of Swaps	Avg Swap Mins	Max Swap Mins	Max Flights Passed	Max Swaps in Hour	Bucket Swaps B2-B1	B3-B2	B4-B3	B5-B4	B6-B5	B7-B6	B3-B1	B4-B2	Diff AC Type & Fix
ALL	1226	137	-4.21	18	25	17	76	45	7	4	2	1	2	0	10
UAL Only	560	47	-4.49	16	11	8	30	13	1	2	0	0	1	0	34
AAL Only	454	38	-4.42	17	9	6	25	10	1	0	1	0	1	0	31
COA Only	23	0	0.00	0	0	0	0	0	0	0	0	0	0	0	0
DAL Only	26	1	-8.00	8	0	1	1	0	0	0	0	0	0	0	0
NWA Only	26	0	0.00	0	0	0	0	0	0	0	0	0	0	0	0
USA Only	22	0	0.00	0	0	0	0	0	0	0	0	0	0	0	0
SLOT Swaps	1226	35	-4.71	-1	17	5	21	11	2	1	0	0	0	0	

Table 5-8: Bucket Swaps at ORD, 4/19/99

5.6 Controlled Time of Arrival (CTA) Restoration in a GDP

Metric Percent of flights in a GDP that can be restored to their FAA-issued controlled time of arrival
Data: CDM (ETMS embellished) flight data
AATT Category: Access

In Reference [7], Metron analyzed the potential for using CTAS to restore flights to their controlled times of arrival (CTAs) as assigned in a ground delay program. The idea behind the restoration was to enforce the concepts of equitable distribution of arrival slots that had been formed by ground delay program procedures. (Hence, the categorization under Access benefit.) Since the actual (attempted) control mechanism in a GDP is a controlled time of departure rather than arrival, flights often arrive outside of their designated arrival time windows for various reasons. Two analyses were performed to assess the potential for incorporation of airline objectives into terminal area scheduling via the integration of CTAS and the Collaborative Decision Making (CDM) program.

Such analyses presume the existence and acceptance of a CAP tool for CTA restoration. At this time, it is not known how feasible or desirable this will be, since it will necessarily involve the speed reduction of flights that are already delayed significantly.

5.6.1 An Upper Bound on CTA Restoration by CTAS

We established a time-based boundary, T, to represent the outer realm of operation of CTAS. This is a radius about the airport of T travel minutes and represents the earliest point at which a flight could be controlled by CTAS. We set the parameter T to the values 45, 60, 90 and 120 to represent a range of (time-based) distances out from the airport that CTAS would first be able to affect a flight.

The delay metric for determining how far a flight f is from its CTA when it crosses the CTAS boundary is determined by

$$delay_T = ETA_T - CTA$$

where

ETA_T = estimated time of arrival of f, T minutes before arrival

CTA = time of last arrival slot assigned to f during the GDP

The assumption here is that ETA_T is the time at which the aircraft would arrive, if allowed to arrive at the airport unimpeded. All ETAs, CTAs and arrival times were recorded by the CDM version of the ETMS operational string (CDM/ETMS). CDM/ETMS rounded these times to the nearest minute.

To model the adjustment of arrival times that would take place by CTAS, we assumed that as each flight *f* crosses the CTAS boundary, CTAS would set a scheduled time of arrival (STA) for *f* as close to its CTA as possible, by adjusting the ETA_T of *f* by *ADJ* minutes, where $\delta_l \leq \text{ADJ} \leq \delta_u$, for some arrival time adjustment parameters $\delta_l > 0$ and $\delta_u > 0$. That is,

$$STA = ETA_T + ADJ$$

where

$$\text{ADJ} = \max(-\delta_l, , delay_T), \text{ if } delay_T \geq 0 , \min(\delta_u,, delay_T) \text{ otherwise.}$$

and

δ_l = max number of minutes a flight can be sped up;

δ_l = max number of minutes a flight can be slowed down.

Since we wanted to explore a range of adjustment parameters and yet keep the number of data sets manageable, we set δ_l and δ_u equal to a single value, δ, which ranged over the values 5, 10, 15 and 20 minutes. In practice, there would probably be more leeway in slowing an aircraft down than in speeding it up. Also, a further limitation would be placed on speed adjustment by the proximity of a flight to other flights (minimum separation). Hence, these results serve as an upper bound on the restoration of flights to their CTAs.

We ran our model of CTAS arrival time adjustment for all flights involved in 13 GDPs implemented at Philadelphia Airport (PHL). The GDPs were selected at random from dates between October 1998 and April 1999. In order to model a range of CTAS boundaries, we set T = 45, 60, 90, and 120. The criteria for flight selection was as follows:

Each flight must have

1. no positive cancellation indicators recorded by CDM/ETMS;
2. a GDP-issued controlled time of arrival (CTA) prior to take-off;
3. a CDM/ETMS-recorded actual runway time of arrival (ARTA);
4. a CDM/ETMS-recorded actual runway time of departure (ARTD)[1] taken at least T minutes prior to its arrival.

---

[1] ARTDs and ARTAs within the CDM/ETMS database are based on ETMS departure (DZ) and arrival (AZ) messages, respectively.

This last criteria (4) implies that a flight may be included in an analysis for some values of T but not others, and that short-haul flights (en-route time less than 45 minutes) were excluded from all of the analyses. Since not every flight will have an ETA recorded exactly T minutes prior to its arrival, ETA_T will not exist for every flight. So, for each boundary value (T = 45, 60, 90, 120) and for each valid flight, we scrolled its CDM-ETMS records for the ETA with the latest time stamp no later than $ARTA - T$.

For each flight k, we computed its scheduled arrival time (STA). (This is the scheduled arrival time as set by CTAS, and should not be confused with the scheduled arrival time set by air carrier operations.) We considered the CTA of the flight to be recoverable if and only if $STA - CTA = 0$. Both ETAs and CTAs are rounded to the nearest minute (by CDM/ETMS). So, this metric allows for one minute of leeway in either direction. These criteria may seem inappropriately strict for some strategic aviation purposes, but the parameters δ_l and δ_u are sufficiently small that a more lenient metric would make little sense.

Our results show that the upper bound on the potential for CTA recovery is quite large. The most conservative result is obtained at the smallest CTAS boundary of T = 45 minutes, and the narrowest window of adjustment, δ =5 minutes. At these values, an average of 39.94% of the CTAs were recoverable (see Table 5-9). As the CTAS boundary grows from T = 45 minutes to T = 120 minutes (moving across the row marked 5 minutes), this figure increases to 48.62%. A much greater CTA recoverability of 77.20% is obtained as the maximum arrival time adjustment δ is relaxed from 5 minutes to 20 minutes (moving down the column marked 45 minutes), with T = 45 fixed. At the other extreme, when both δ and T are at their maximum (20 minutes and 120 minutes, respectively), 84.38% of CTAs are recoverable.

Since very few of the flights were on time at the CTAS boundary (i.e., $ETA_T = CTA$), virtually all flights (93.92 to 96.51 percent) would be subject to an arrival time adjustment by CTAS. The average amount of absolute adjustment, or 'work', that would be done per flight toward CTA restoration would be 4.04 minutes and 8.06 minutes, at the two extremes, (δ, T) = (5, 45) and (20, 120), respectively. See Table 5-10. For all settings of (δ, T), the average amount of arrival time adjustment is slightly negative, indicating that, on average, flights would be sped up by CTAS. This is to be expected based on the table in Table 5-11, which gives a breakdown of CTA delays prior to adjustment. We see that those flights off their CTA when entering the CTAS realm are almost evenly split between those early and those late but the absolute delay is much greater for those running late. For instance, when T = 60, the average delay for early flights is 5.24 while the average delay for late flights is 11.32 minutes. This leads to an average delay of 6.08 minutes and skews the average adjustment in the negative direction.

To see the type of effect that CTAS adjustments would have on the distribution of CTA compliance, consider the two pie charts in Table 5-12. For flights involved in a GDP at

PHL on 04 March 1999, $T = 60$, $\delta = 5$, these charts show the percentages of flights that were early, late, and on-time before and after CTAS adjustments. Notice that the percentage of (CTA) on-time flights rose dramatically from 11% to 47%, for an increase of over 327% in the on-time category. The percentage of late flights dropped from 47% to 35%, for a categorical decrease of over 25%, while the percentage of early flights dropped from 42% to 18%, for a categorical decrease of over 57%.

In order to condense our results into a more memorable format, we selected a CTAS boundary value and a maximum adjustment value that would most likely be encountered in practice. This gave us a single upper bound on CTA recoverability that is, in some sense, typical.

We selected for $T = 60$ for the CTAS boundary, since this is the approximate flight time from an ARTCC boundary to an airport located in the middle of the ARTCC. It is less clear what to select for the flight adjustment range, since this varies with aircraft type, velocity, and several other parameters. On average, this maximum adjustment parameter probably lies somewhere between $\delta = 5$ and $\Delta = 10$ minutes. By examining our results in Table 5-9 corresponding to $T = 60$, we interpolated between $\delta = 5$ and $\delta =10$ to find that the upper bound on CTA recovery is approximately 50% of the flights, with an average arrival time adjustment of about 5 minutes. At these parameters, approximately 95% of the flights would have their arrival time adjusted by at least one minute.

5.6.2 Best Estimate of CTA Recoverability

The purpose of this analysis is to address as many of the simplistic assumptions in Section 2.2 as possible and, by so doing, produce a better estimate of the potential for CTA restoration. There are three basic limitations to arrival time scheduling by CTAS that must be addressed. First is that arrival time adjustment is limited by the physical amount an aircraft can be sped up or slowed down in the limited time it has before final approach. We set the arrival adjustment parameters to the values $\delta_1 = 5$ and $\delta_u =10$, as we felt these would be the values most likely be encountered in practice.

The second limitation is that the arrival time of a flight can be adjusted without consideration to the positioning of other aircraft. Flights must maintain a minimum separation and, though there are exceptions, flights within a single arrival fix stream generally cannot overtake each other. To model arrival fix streams, we partitioned the flights according to the arrival fixes that they actually used (historical hindsight from CDM/ETMS data) and formed a queue for each fix, with order being determined by their ETAs taken at the CTAS boundary, i.e., about 60 minutes away from arrival. Since we maintained order within this queue, this analysis serves as a lower bound for CTA restoration.

Lastly, we insisted that arrival time adjustments be made in accordance with the acceptance rate of the airport. Through CAP-FSM interfacing, we are assuming CTAS knowledge of GDP parameters, one of which is the planned airport acceptance rate (PAAR). PAARs are set well in advance of a GDP based on the conversion of weather forecasts to runway configurations then to historical acceptance rates. Due to the difficulty of forecasting weather several hours in advance, the actual acceptance rate at the airport upon which CTAS bases its schedule may not be the same as the PAAR. But, for the purposes of this analysis, we assume that CTAS schedules arrival acceptance rates in accordance with the PAARs.

In order to avoid generating a plethora of results, we selected a single value for the CTAS boundary, T = 60 minutes; we felt this is the most realistic approximation of the (hypothetical) CTAS operational radius around an airport. In summary, we simulated constraints on CTAS arrival time adjustments as follows:

1. each arrival time adjustment must fall in the range [–5, 10];
2. no flight can overtake another flight within its arrival fix queue;
3. airport capacity is equal to the AAR in CTAS (which we assume to be the PAAR in the GDP).

This last constraint dictates that minimum spacing between flights is determined by apportioning to each arrival fix a capacity in proportion to its demand. For instance, if there are two arrival fixes with respective traffic demands of 40% and 60%, and if the planned airport acceptance rate (PAAR) for a given hour is 35 flights, then the arrival streams are allotted 35(.40) = 14 and 35(.60) = 21 landing slots respectively, uniformly distributed over that hour, with corresponding minimum spacing. This minimum spacing in our model may be violated: given the [-5, 10] restriction on flight movement, there could be sufficient demand such that the AAR is exceeded for a given hour (CTAS would not actually allow this to happen).

The flow of traffic over each arrival fix was modeled as a queue, with order of arrival determined by the ETA taken 60 minutes prior to runway time of arrival (historical hindsight). The arrival time adjustment ADJ_k for the k^{th} flight was computed as follows. Let

ETA_k = estimated time of arrival of flight k

STA_k = CTAS scheduled time of arrival of flight k

CTA_k = GDP-controlled time of arrival of flight k

$Buff$ = desired buffer (separation) between successive flights

$Dist_k$ = $ETA_k - STA_{k-1}$, if $k > 1$, infinity, otherwise.

$Fspace_k = \max(Dist_k - Buff, 0)$

$\delta_k = ETA_k - CTA_k$

Buff is set in accordance with the AAR, as described earlier; $Dist_k$ is the distance between successive flights prior to arrival time adjustments; $Fspace_k$ is the amount of `free space' between successive flights, that is, the maximum number of minutes that a flight can be advanced without violating any of the constraints. The final scheduled time of arrival of the k^{th} flight (STA_k) is determined by

$$STA_k = ETA_k + ADJ_k .$$

Arrival time adjustments were set by processing the flights in order of increasing ETA. For each flight k, we set

$$ADJ_k = \max(\delta_l, -\Delta_k, -Fspace_k), \text{ if } \Delta_k > 0, \max(\delta_u, -\Delta_k), \text{ otherwise.}$$

To facilitate the flight processing, we considered only negative adjustments (speed ups) for each flight in the range $[\delta_u - \delta_l, 0]$, by first adding δ_u minutes of delay to each flight. That is, for all k, we set $ETA_k = ETA_k + \delta_u$ and

$$ADJ_k = \max(|\delta_u\text{-}\ \delta_u|, -\Delta_k, -Fspace_k), \text{ if } \Delta_k > 0, 0, \text{ otherwise.}$$

The final adjustment is obtained by the translation $ADJ_k = ADJ_k + \delta_u$. We set the lower and upper bounds on ADJ_k to $\delta_l = -5$ and $\delta_u = 10$, respectively.

We ran our CTAS simulator on all flights with controlled times of arrival at PHL and SFO for 13 randomly selected GDP days. There are four things to be considered in the interpretation of the results: CTA recoverability, the amount of movement (in minutes), position shifting, and the effect on the arrival rate at the airport (AAR). We discuss each one in turn.

5.6.3 CTA Recoverability

Our results show that after the addition of operational and physical constraints to arrival time adjustment, there is still considerable potential for CTA restoration. An average of 21.98% of the flights involved in GDPs at PHL could be restored to their CTA (19.21% at SFO) with a CTAS boundary of 60 minutes and an adjustment range of $\delta_l = -5$ to $\delta_u = 10$. The results for each GDP date appear in Table 5-13 in the column marked "Queue".

In order to make a direct comparison with the upper bound results Section 1, we recomputed the percentages of flights that could be restored to their CTA without any arrival fix modeling or adherence to AAR (that is, using only the bounding parameters $\delta_l = -5$ and $\delta_u = 10$). The results are in the column marked "No queue" of Table 5-13.

There is a significant drop in the CTA restoration once the arrival fix modeling is introduced. There are two reasons for this. First, the acceptance rate of the airport (assumed to be equal to the PAAR of the GDP) can only have the effect of pushing flights back (later) in time. Surges in demand (based on ETA) will be smoothed out by this operational constraint. In fact, a flight that is already delayed (off its CTA) could be delayed even further by this constraint. But, the instances in which the PAAR became a factor were rare.[2] A second, more dominant factor is the constraint that no flight within an arrival queue could overtake another flight in its queue. This has the effect of inhibiting forward movement of flights in the queue but not backward. Each flight was processed as it crossed the CTAS boundary, with only knowledge of the position of flights that preceded it. Each of these factors constrains the forward movement (speeding up) of flights. Greater CTA restoration is likely to take place when flights are ahead of their CTA.

It is dubious that CTAS would leave such large buffers in the arrival streams over fixes. So, we consider the values in the column marked "Queue" to be a lower bound, while the values in the "No Queue" column are an upper bound. Based on the averages over all days, our conclusion is that the percentage of flights at PHL that could be restored to their controlled time of arrival (CTA) by CTAS lies somewhere between 44.73% and 21.98% (40.27% and 19.21% at SFO), with the most likely estimate lying closer to the latter figure(s).

5.6.4 The Amount of Movement

The average arrival time adjustment (over all arrival fixes) at PHL was 5.62 minutes with an average absolute adjustment of 6.84 minutes (5.72 and 6.99 minutes, respectively, for SFO). See Table 5-14. The values were fairly uniformly distributed over the arrival fixes. This is a significant shift in the positive direction (adding delay to flights) from our results in the prior analysis: for the boundary value $T = 60$, the average adjustment was in the range −0.75 to −2.84, meaning that on average flights were sped up slightly. As mentioned in Section 3, the combination of airport capacity adherence and order preservation over arrival fixes tends to dampen the ability of CTAS to speed flights up to meet their CTA. Also, in the prior results (Table 5-9), the bounding parameters were set equal but in this analysis, more allowance has been made for slowing flights down than speeding them up. All of these factors combine to skew the distribution of arrival time adjustments toward the positive. The average absolute adjustment (amount of 'work') done under this analysis (6.84 at PHL) was in keeping with the results of the prior analysis (a range from 4.04 to 8.06 minutes).

[2] Some flights were disqualified from the study due to corrupted or missing records, which lowers the demand that we input to our CTAS simulator from what it actually may have been.

5.6.5 Position Shifting

Although our CTAS simulator maintains arrival order (by ETA) within each arrival fix stream, this does not prevent re-ordering within the overall arrival stream. Another valid measure of arrival time movement is position shift within an arrival queue. In our analysis, we arranged flights by increasing ETA (taken 60 minutes out from arrival) as they crossed the CTAS boundary of operation to recreate the order in which flights would be processed by CTAS. The flights were partitioned into arrival stream classes, each with its own arrival order. We checked for position shifts by a method similar to the fast-time simulation done by Carr et al. in their analysis of banking priorities within CTAS scheduling [1].

Table 5-15 shows that the largest position shifts were in the negative direction, that is, moving a flight to an earlier position in the queue. The minimum recorded shift at PHL was −10 positions while the maximum was 6 (−9 and 5, resp., for SFO). The average minimum and maximum position shifts ranged from −8.23 to 4.31 at PHL (−5.46 to 3.08 at SFO). For any given data set, the position shifts average to zero, hence, are not recorded. The average absolute position shifts are generally small positive values (averaging 1.47 at PHL and 3.08 at SFO); combined with the ranges of the distributions, this indicates that the bulk of the flights were shifted to later positions in the queue (delayed).

Percentage of Flights with Recoverable CTA - PHL				
Max Adjust δ	45 min out	60 min out	90 min out	120 min out
5 min	39.94	42.32	47.48	48. 2
10 min	60.41	62.90	66.74	68.5
15 min	71.51	73.83	77.19	78.
20 min	77.20	78.31	81.94	84.3

Table 5-9: CTA recoverability at PHL

Feasible CTA Adjustments – PHL				
Adjustment	45 min out	60 min out	90 min out	120 min out
Avg 05	-0.59	-0.75	-0.68	-0.33
Avg 10	-1.38	-1.55	-1.24	-0.69
Avg 15	-2.08	-2.28	-1.77	-1.02
Avg 20	-2.65	-2.84	-2.17	-1.33
Avg Abs 05	4.04	3.94	3.78	3.69
Avg Abs 10	6.60	6.37	5.99	5.77
Avg Abs 15	8.33	7.99	7.40	7.11
Avg Abs 20	9.61	9.21	8.44	8.06

Table 5-10: Feasible CTA adjustments

CTA Delays - PHL				
	45 min out	60 min out	90 min out	120 min out
Avg Numb of Flights	147.36	121.71	76.07	49.
Percent Early	43.36	40.85	41.31	44. 1
Percent Late	53.15	54.73	53.85	49. 1
Percent On-time	3.49	4.41	4.84	6.
Avg Delay	6.07	6.08	4.58	2. 1
Avg Abs Delay	17.56	16.57	13.67	12.1
Avg Delay (Early)	-5.75	-5.24	-4.55	-5.
Avg Delay (Late)	11.81	11.32	9.12	7.1

Table 5-11: Breakdown of CTA delays

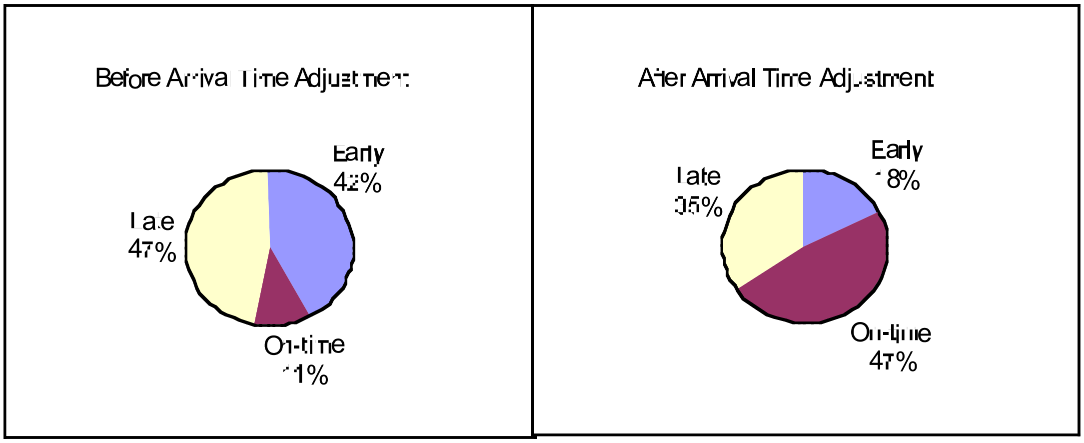

Table 5-12: Distribution of delay (ETA – CTA) before and after adjustment

Percent recoverable Flights [-5, 10]					
PHL			SFO		
Date	No queue	Queue	Date	No queue	Queue
01/22/99	42.91	25.26	(1) 01/07/99	43.24	[illegible]
02/07/99	44.81	15.85	(2) 01/07/99	1.06	[illegible].35
02/28/99	50.85	20.17	01/15/99	51.75	25.17
03/03/99	43.95	17.20	01/16/99	45.95	27.[illegible]3
03/04/99	47.39	20.92	01/17/99	56.45	22.2
03/14/99	32.43	18.92	01/18/99	4.00	12.[illegible]
03/21/99	37.16	20.27	01/19/99	53.74	24.[illegible]
04/01/99	51.88	28.13	01/20/99	50.00	25.25
04/16/99	47.67	23.26	01/22/99	52.14	22.75
10/08/98	47.58	23.39	01/23/99	44.57	16.57
10/09/98	36.43	18.22	01/26/99	46.76	14.3
11/20/98	55.29	30.59	01/30/99	38.46	23.[illegible]
12/29/98	43.14	23.53	01/31/99	35.47	33.33
Average	**44.73**	**21.98**	**Average**	**40.27**	**19.21**

Table 5-13: CTA recoverability at PHL and SFO, with and without arrival fix modeling

Adjustments by Arrival Fix							
PHL				SFO			
AFIX	Avg	Avg Abs	% Traffic	AFIX	Avg	Avg Abs	% Traffic
BUNTS	5.51	6.50	35.81	BRINY	7.25	7.83	20.18
DQO	6.31	7.30	9.42	CEDES	5.54	6.74	48.35
MAZIE	4.36	6.49	6.38	PYE	4.92	6.68	25.48
PTW	6.73	7.43	6.47	SKUNK	5.17	6.70	18.61
TERRI	5.63	6.96	18.13				
VCN	5.17	6.39	23.79				

Table 5-14: Average adjustments (in minutes) by arrival fix at PHL and SFO

Arrival Order Position Shifts							
PHL				SFO			
Date	MIN	MAX	Avg Abs	DATE	MIN	MAX	Avg Abs
01/22/99	-9	5	1.56	(1)1/7/99	-5	2	0.3
02/07/99	-10	4	1.22	(2)1/7/99	-2	1	0.11
02/28/99	-9	6	1.61	01/15/99	-7	4	1.0
03/03/99	-10	4	1.40	01/16/99	-5	4	0.83
03/04/99	-10	5	1.87	01/17/99	-5	4	0.8
03/14/99	-7	6	1.28	01/18/99	-2	1	0.1
03/21/99	-7	4	1.28	01/19/99	-7	5	1.7
04/01/99	-8	4	1.89	01/20/99	-7	3	0.92
04/16/99	-5	3	1.33	01/22/99	-5	4	0.85
10/08/98	-7	5	1.37	01/23/99	-9	3	0.9
10/09/98	-8	4	1.85	01/16/99	-5	4	0.75
11/20/98	-6	3	1.15	01/30/99	-6	2	1.0
12/29/98	-11	3	1.35	01/31/99	-6	3	1.0
Average	**-8.23**	**4.31**	**1.47**	**Average**	**-5.46**	**3.08**	**0.81**

Table 5-15: Position shifts in arrival order at PHL and SFO

6 Metric and Benefits Summary Tables

Table 6-1 summarizes each decision support tool (DST) function of CAP, the corresponding benefits that are expected to be derived, the metric employed in this study, the AATT high-level category under which the metric falls, and any economic conversions (if applicable). Benefits are hypothetical for the year 1999, based on the top 38 airports in the NAS for the year 1999 (see Table 8-1).

Table 6-2 is similar to Table7-1 only the derived benefits are for the year 2015, based on the top 38 airports in the NAS for the year 2015 (see Table 6-1). Although the total number of arrivals is predicted to rise considerably by 2015, we do not anticipate any change in the benefits computations that are based on a percentage of total arrivals. Therefore, the percentage-based figures in Table 6-1 (e.g., percent increase in on-time arrivals) have been carried over to Table 6-2. However, the figures in numbers of flights or dollars would increase at the same rate of traffic growth. Using the total numbers of arrivals in Table 6-1, we computed the increase of arrivals that would take place from 1999 to 2015:

$$\text{Ratio} = (\text{Number of arrivals in 2015}) / (\text{Number of arrivals in 1999}) = 12{,}048{,}417 / 8{,}448{,}220 \approx 1.43$$

This ratio was used to increase the number of arrivals NAS-wide at the 38 top airports or the number of arrivals at a hubbing operation. For instance, we applied the ratio to the number of flights in cell F9 of Table 6-1 (2,855) to obtain 4,071 flights in cell F9 of Table 6-2:

$$(\text{Number of flights in 1999}) \times (\text{Percentage increase from 1999 to 2015}) = 2{,}855 \times 1.43 = 4{,}071$$

The ratio was also used to compute the dollar figures in Table 6-2 based on the dollar figures in Table 6-1. Since no adjustment has been made for inflation, all dollar figures in Table 6-2 are in 1999 dollars.

Table 6-1: NAS-wide Benefits (Hypothetical) for Year 1999

	A	B	C	D	E	F	G	H
1	Report Section	CAP DST Function	Anticipated Benefit	AATT Category	Metric	Lower Bound	Upper Bound	Best Estimate
2	5.1	Passive transmission of CTAS data to the air carriers	Improved ETA Accuracy Provided by CTAS	Predict	Weighted average of ETA errors (ETA – actual arrival time)	30.28 seconds	64.00 seconds	47.76 seconds
3	5.2 5.3		Improved Hold-go and diversion decisions based on increased ETA accuracy.	Predict	Instances in which only the CTAS ETA (or ETMS ETA) fell within range of arrival time	Pending further analysis	Pending further analysis	Pending further analysis
4	5.4		Improved prediction of release time from a state of airborne holding	Predict	Same as 5.1 only restricted to flights that experienced airborne holding	Pending further analysis	Pending further analysis	Pending further analysis
5	6.3	*Inter*-carrier Delay Exchange Program conducted through CAP	Increased on-time arrival performance	Flex	Increased on-times	1.81%	5.77%	1.81%
6					Increased on-times	152,931 Flts	487,520 Flts	152,931 Flts
7					Delay savings	$29,914,527	$95,362,812	$29,914,527
8	6.4	*Intra*-carrier Delay Exchange Program conducted through CAP	Increased on-time arrival performance per hubbing operation	Flex	Increased on-times	2.01%	5.05%	2.01%
9					Increased on-times	2,855 Flts	7,172 Flts	2,855 Flts
10					Delay savings	$558,339	$1,492,991	$558,339
11	6.5	Arrival Time Adjustments	Mitigate effects of delays	Flex	Increased on-times	2.06%	2.06%	2.06%
12					Increased on-times	174,054 Flts	174,054 Flts	174,054 Flts
13					Delay Savings	$21,740,090	$21,740,090	$21,740,090
14	6.6	CTA Restoration through CAP-FSM	Restore GDP equity	Access	Percentage of flights restored to CTA	19.21%	44.73%	19.21%

Note: Dollar figures are in 1999 dollars.

Table 6-2: NAS-wide Benefits (Hypothetical) for Year 2015								
	A	B	C	D	E	F	G	H
1	**Report Section**	**CAP DST Function**	**Anticipated Benefit**	**AATT Category**	**Metric**	**Lower Bound**	**Upper Bound**	**Best Estimate**
2	5.1	Passive transmission of CTAS data to the air carriers	Improved ETA Accuracy Provided by CTAS	Predict	Weighted average of ETA errors (ETA – actual arrival time)	30.28 seconds	64.00 seconds	47.76 seconds
3	5.2 5.3		Improved Hold-go and diversion decisions based on increased ETA accuracy.	Predict	Instances in which only the CTAS ETA (or ETMS ETA) fell within range of arrival time	Pending further analysis	Pending further analysis	Pending further analysis
4	5.4		Improved prediction of release time from a state of airborne holding	Predict	Same as 5.1 only restricted to flights that experienced airborne holding	Pending further analysis	Pending further analysis	Pending further analysis
5	6.3	*Inter*-carrier Delay Exchange Program conducted through CAP	Increased on-time arrival performance	Flex	Increased on-times	1.81%	5.77%	1.81%
6					Increased on-times	218,077 Flts	695,194 Flts	218,077 Flts
7					Delay savings	$42,657,511	$135,985,443	$42,657,511
8	6.4	*Intra*-carrier Delay Exchange Program conducted through CAP	Increased on-time arrival performance per hubbing operation	Flex	Increased on-times	2.01%	5.05%	2.01%
9					Increased on-times	4,071 Flts	10,227 Flts	4,071 Flts
10					Delay savings	$796,266	$2,000,568	$796,266
11	6.5	Arrival Time Adjustments	Mitigate effects of delays	Flex	Increased on-times	2.06%	2.06%	2.06%
12					Increased on-times	248,197 Flts	248,197 Flts	248,197 Flts
13					Delay Savings	$31,000,929	$31,000,929	$31,000,929
14	6.6	CTA Restoration through CAP-FSM	Restore GDP equity	Access	Percentage of flights restored to CTA	19.21%	44.73%	19.21%

Note: Dollar figures are in 1999 dollars.

7 Future Recommendations and Analyses

In this study, we made several large assumptions about CAP functionality. Future studies of the benefits associated with a CAP-like tool would be helped most by a more refined development of CAP functions. In addition, we have listed below other recommendations for future studies.

Section 4.2: We chose a parameter 15 minutes was chosen as the minimum time that would be required at the hubbing airport for passengers to make their connecting flights. Through informal experimentation, we found that our results were relatively invariant with the value of this parameter.) This experimentation could be made more formal. Also, surveys could be taken to see how far in advance of departure AOCs tend to make critical hold/go or diversion decisions.

Section 5.3:
In our analysis of delay bucket exchanges, we reported mainly the results from movements into the on-time delay bucket. This was largely driven by the inability to associate financial costs with the other six delay buckets and the commitment to avoiding the use of proprietary data. Perhaps a study could be done of airline costs for the delay buckets we have cited.

We computed the cost savings from movement of flights from delay bucket B_2 to delay bucket B_1 by multiplying the total number of such movements by the cost difference between the two buckets. If a cost in dollars C_n can be associated with each delay bucket B_n, then the total cost of all bucket movements could be computed as follows:

The results will vary with the traffic demand and airport capacity. Higher demand or lower capacity will induce more CTAS imposed delay and allow a greater number of delay exchanges.

Section 5.3, 5.4, 5.5: In Section 5.3 we applied or removed two restrictions on delay exchanges to obtain lower and upper bounds, respectively, on benefits. One of these was the no-passing restriction, which, in essence, said that a flight could not overtake another flight it the same arrival stream class. Better data might be available for this.

In Section 5.5, we studied the benefits of arrival time adjustments, which can be viewed as a form of pure delay exchange, that is, movement of flights into vacant arrival slots. This was intentionally isolated from the benefits of Section 5.3 and 5.4, in which flights were moved into slot occupied by other flights. In practice, a delay exchange program would probably be a blend of these two approaches. Further work could be done to assess the value of such a program.

8 Supporting Data

Table 8-1: Top 38 Airports (by Operations) in 1999 or 2015					
		A	B	C	D
Row	Airport	Ops 1999	Ops 2015	Rank-1999	Rank-2015
1	ANC	307792	450613	34	31
2	ATL	879386	1314590	2	2
3	BOS	516366	587781	8	21
4	BWI	307368	415555	35	33
5	CLE	327921	501571	32	28
6	CLT	451768	556207	19	26
7	CVG	459053	739434	16	10
8	DCA	330138	351454	31	40
9	DEN	498285	721878	12	12
10	DFW	875400	1320511	3	1
11	DTW	542343	822085	6	7
12	EWR	462501	606433	15	18
13	HNL	344250	495375	30	29
14	IAD	438260	576812	20	22
15	IAH	458741	757991	17	9
16	IND	250682	377379	44	38
17	JFK	355461	414149	29	34
18	LAS	506215	851273	10	6
19	LAX	762828	1171525	4	3
20	LGA	365178	412918	28	35
21	MCO	365339	614998	27	17
22	MDW	289145	385244	37	36
23	MEM	367155	568441	26	24
24	MIA	530284	735483	7	11
25	MSP	497089	782380	13	8
26	OAK	507170	605593	9	19
27	ORD	899325	1151324	1	4
28	PDX	327664	491527	33	30
29	PHL	479144	700800	14	14
30	PHX	545520	903341	5	5
31	PIT	458162	568695	18	23
32	SAT	264873	381414	41	37
33	SEA	411332	561758	23	25
34	SFB	380777	709342	24	13
35	SFO	434632	656029	21	16
36	SJC	304168	439836	36	32
37	SLC	369579	599832	25	20
38	SNA	433867	530774	22	27
39	STL	503324	689145	11	15
40	TUS	277133	340547	38	42
41	TotalOps	16898440	24096834		
42	TotalArrv	8449220	12048417		

Table 8-2: GDP Days in Year 1999			
	A	B	C
Airport	GDP days	Non-GDP Days	Total
ORD	80	285	365
PHL	43	322	365
DFW	3	362	365

9 References

[1] Heere, Karen R., and Richard E. Zelenka, Richard E., “A Comparison of Center/TRACON Automation System and Airline Time of Arrival Predictions”, NASA Research Document NASA/TM-2000-209584, February, 2000.

[2] Zelenka, Richard E., Cheryl Quinn, Roger Beatty, and Karen Heere, “The Impact of CTAS Information on Airline Operational Control”, Air Traffic Control Quarterly, Vol. 8(1) 33-62 2000.

[3] Ball, M. O., et al. “Collaborative Decision Making in Traffic Flow Management: A Preliminary Assessment”, NEXTOR report #RR-99-3 to the FAA, 1998.

[4] FAA FFP1 Report, “An Operational Assessment of Collaborative Decision Making in Air Traffic Management: Measuring User Impacts through Performance Metrics”, Document Control Number: R90145-01, December 1999.

[5] Hoffman, Robert L., “What is IPE?”, the NEXTOR Web Site, http//www.Nextor.umd.edu, October 1998.

[6] Hoffman, Robert L., “The Integrated Predictive Error (IPE) Metric”, manuscript in preparation.

[7] Hoffman, R., Chris Brinton, Noelle Shaw, Gregory Carr, “The Potential for Arrival Resequencing Within the CTAS Environment”, manuscript in preparation.

[8] Hoffman, R., Chris Brinton, Noelle Shaw, “Concept Exploration of CTAs-CAP/CDM-FSM Integration”, a study conducted by Metron for NASA Ames Research Center.

[9] Carr, G. C., Erzberger, H., Neuman, F., “Airline Prioritization in Sequencing and Scheduling”, 2nd USA/Europe Air Traffic Management R&D Seminar, Orlando, FL, December 1998.

[10] Hoffman, R. “Integer Programming Models for Ground Holding in Air Traffic Flow Management”, University of Maryland Ph.D. dissertation, 1999.

[11] *Economic Values for Evaluation of Federal Aviation Administration Investment and Regulatory Programs*. Prepared by the Federal Aviation Administration’s Office of Aviation Policy and Plans (APO).