PB2002-104359

Project Report
NASA/A-5

Contributions to the AIAA Guidance, Navigation & Control Conference

S.D. Campbell
Editor

23 January 2002

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS

Prepared for the National Aeronautics and Space Administration
Ames Research Center, Moffett Field, CA 94035
and the Federal Aviation Administration
Washington, DC 20591

This document is available to the public through
the National Technical Information Service,
Springfield, Virginia 22161.

TECHNICAL REPORT STANDARD TITLE PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No
NASA/A-5		

4. Title and Subtitle	5. Report Date
Contributions to the AIAA Guidance, Navigation & Control Conference	23 January 2002
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Steven D. Campbell, Editor	NASA/A-5

9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02420-9108	11. Contract or Grant No. NASA Ames

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
NASA Ames Research Center Moffett Field, CA 94035 Federal Aviation Administration Washington, DC 20591	Project Report
	14. Sponsoring Agency Code

15. Supplementary Notes

This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology, under Air Force Contract F19628-00-C-0002.

16. Abstract

This report contains six papers presented by the Lincoln Laboratory Air Traffic Control Systems Group at the American Institute of Aeronautics & Astronautics (AIAA) Guidance, Navigation and Control (GNC) conference on 6-9 August 2001 in Montreal, Canada. The work reported was sponsored by the NASA Advanced Air Transportation Technologies (AATT) program and the FAA Free Flight Phase 1 (FFP1) program. The papers are based on studies completed at Lincoln Laboratory in collaboration with staff at NASA Ames Research Center.

These papers were presented in the Air Traffic Automation Session of the conference and fall into three major areas: Traffic Analysis & Benefits Studies, Weather/Automation Integration, and Surface Surveillance. In the first area, a paper by Andrews & Robinson presents an analysis of the efficiency of runway operations at Dallas/Ft. Worth using a tool called PARO, and a paper by Welch, Andrews, & Robinson presents delay benefit results for the Final Approach Spacing Tool (FAST). In the second area, a paper by Campbell, et al. describes a new weather distribution system for the Center/TRACON Automation System (CTAS) that allows ingestion of multiple weather sources, and a paper by Vandevenne, Lloyd, & Hogaboom describes the use of the NOAA Eta model as a backup wind data source for CTAS. Also in this area, a paper by Murphy & Campbell presents initial steps towards integrating weather-impacted routes into FAST. In the third area, a paper by Welch, Bussolari, and Atkins presents an initial operational concept for using surface surveillance to reduce taxi delays.

17. Key Words	18. Distribution Statement
	This document is available to the public through the National Technical Information Service, Springfield, VA 22161.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	80	

FORM DOT F 1700.7 (8-72) Reproduction of completed page authorized

ABSTRACT

This report contains six papers presented by the Lincoln Laboratory Air Traffic Control Systems Group at the American Institute of Aeronautics & Astronautics (AIAA) Guidance, Navigation and Control (GNC) conference on 6-9 August 2001 in Montreal, Canada. The work reported was sponsored by the NASA Advanced Air Transportation Technologies (AATT) program and the FAA Free Flight Phase 1 (FFP1) program. The papers are based on studies completed at Lincoln Laboratory in collaboration with staff at NASA Ames Research Center.

These papers were presented in the Air Traffic Automation Session of the conference and fall into three major areas: Traffic Analysis & Benefits Studies, Weather/Automation Integration and Surface Surveillance. In the first area, a paper by Andrews & Robinson presents an analysis of the efficiency of runway operations at Dallas/Ft. Worth using a tool called PARO, and a paper by Welch, Andrews & Robinson presents delay benefit results for the Final Approach Spacing Tool (FAST). In the second area, a paper by Campbell, *et al* describes a new weather distribution system for the Center/TRACON Automation System (CTAS) that allows ingestion of multiple weather sources, and a paper by Vandevenne, Lloyd & Hogaboom describes the use of the NOAA Eta model as a backup wind data source for CTAS. Also in this area, a paper by Murphy & Campbell presents initial steps towards integrating weather impacted routes into FAST. In the third area, a paper by Welch, Bussolari and Atkins presents an initial operational concept for using surface surveillance to reduce taxi delays.

TABLE OF CONTENTS

AIAA Guidance, Navigation & Control Conference, Montreal, Quebec, August 6-9, 2001 AIAA-2001-4359

RADAR-BASED ANALYSIS OF THE EFFICIENCY OF RUNWAY USE*

John W. Andrews
MIT Lincoln Laboratory
244 Wood Street
Lexington, MA 02420-9108

John E. Robinson III
NASA Ames Research Center
Moffett Field, CA 94035-1000

ABSTRACT

The air transportation system faces a challenge in accommodating growing air traffic despite an inability to build new runways at most major airports. One approach to alleviating congestion is to find ways of using each available runway to the maximum extent possible without violating safety standards. Some decision support tools, such as the Final Approach Spacing Tool (FAST) that is a part of the Center TRACON Automation System (CTAS), are specifically targeted toward achieving greater runway throughput by reducing the average landing time interval (LTI) between arrivals at a given runway. In order to understand the potential benefits of such innovations, techniques for detecting spacing inefficiencies and estimating potential throughput improvements are needed. This paper demonstrates techniques for analyzing radar data from actual airport operations and using it to validate, calibrate, and extend analyzes of the FAST benefits mechanisms. The emphasis is upon robust statistical measures that can be produced through automated analysis of radar data, thus enabling large amounts of data to be analyzed.

INTRODUCTION

A number of analytic and simulation studies have attempted to assess the potential benefits resulting from deployment of the Final Approach Spacing Tool (FAST) that is a part of the Center TRACON Automation System (CTAS).[1,2,3,4] One of the primary sources of FAST benefits is the increased precision of control, which is presumed to reduce the average landing time intervals (LTIs) at each runway. In general, it is assumed that achieved separations contain some amount of excess spacing (not required by separation standards) and that by allowing more precise control, this excess is reduced in a uniform way for all arrivals to which FAST advisories are applied. By saving a few seconds of runway time for each arriving pair, this mechanism provides an increase in runway capacity. The delay savings that accrue over an extended period of operation are found by integrating the delay reductions achieved over a variety of traffic and weather conditions.

In this paper, data from actual airport operations is analyzed and applied to the problem of validating, calibrating, and extending the model for the key FAST benefits mechanism – landing time interval reduction. The analysis of actual operations data is also helpful in prioritizing research activities to focus upon areas where the greatest opportunity lies. This work extends the capabilities used in earlier data analysis conducted by Boswell and Ballin and Erzberger.[1,2] The emphasis is upon robust statistical measures that can be produced through automated analysis routines, thus enabling large amounts of data to be analyzed.

PACKAGE FOR ANALYSIS OF RUNWAY OPERATIONS (PARO)

All major airports acquire and archive radar data on traffic in the terminal area using the Automated Radar Terminal System (ARTS). When combined with basic flight plan information and knowledge of the runway layout, this data provides insight into the flow rates into the terminal and the manner in which particular runways were being utilized. A software package called Package for Analysis of Runway Operations (PARO) was written to automatically process such data and produce analyses relevant to the efficiency of operations. PARO is written in the C++ programming language. Data analyses presented in this paper will focus primarily upon analysis of four DFW data sets that were available during the software development period. These data sets were used to develop the analysis techniques and give some preliminary insight into AFAST benefits questions. Analysis of additional sets of data are being analyzed currently.

PARO processing takes place in three phases designated G0, G1 and G2. Phase G0 involves the reading of raw data files, correcting certain errors and anomalies, and producing new radar data files. In the original data files, tracks appear in order of the time of the first radar report in the track. Phase G1 involves reading the G0 radar data files, correcting and

*This work was performed for the National Aeronautics and Space Administration under Air Force Contract No. F19628-00-C-0002.

validating the input data, estimating velocities, and conducting certain analyses that require complete track data. Phase G2 involves reading and processing the summary data files produced from G1 processing. Among the variables that may be analyzed are the path length flown, the time of crossing the outer marker, the interarrival interval relative to the preceding arrival, etc. By operating only upon the summary files, G2 analyses can run more rapidly without having to process the more voluminous track data.

Bayesian Runway Assignment Procedure (BRAP)

The ability to properly assign each observed operation to a particular runway is essential for reliable analysis of multi-runway operations. If radar data were complete and of sufficient accuracy, such assignment might require a simple comparison of the surface intercept projection of tracks with the known runway locations. However, several imperfections in the radar data (particularly altitude coverage limitations) lead to the need for a somewhat more sophisticated approach to runway assignment.

A Bayesian approach to runway assignment has been developed as part of PARO. Under the Bayesian approach, the runway assignment is viewed both as a parameter that determines the likelihood of any given set of radar observations and as a random variable that has its own probability distribution. The Bayesian approach allows an optimum utilization of all available information about how runways are being used and what was observed with radar. The result is a runway assignment algorithm that is more accurate than any algorithm based solely upon radar data for a single track.

Data Completeness

The completeness of the data is of great concern when evaluating the efficiency of airport operations. Missing tracks create gaps in the arrival stream that can be mistakenly attributed to system inefficiencies. As a general rule, data should be approximately 99% complete to perform all the PARO analyses of interest. (That is, not more than 1 aircraft in 100 should be missing from the set of radar tracks). The DFW data is judged to be adequate in this respect.

GENERAL INSPECTION OF AIRPORT OPERATIONS AT DFW

In this section we will discuss some general attributes of the traffic flow that are relevant to the analysis of interarrival spacing. Figure 1 shows the runway layout at DFW. There are seven runways and thus 14 possible landing directions. When traffic is flowing to the north, the airport is said to be in a "north flow". When traffic is flowing to the south, the airport is said to be in a "south flow".

Figure 2 shows a selection of tracks plotted during a period when the airport was in a north flow. It is difficult to determine exactly how efficiently the runways were being used by casual inspection of the actual tracks. However, the plots and analyses that will now be described are designed to provide insight into this question.

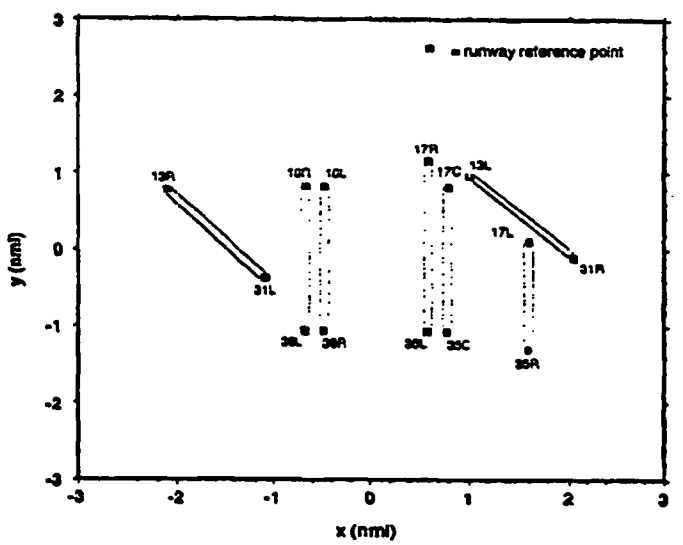

Figure 1. Runway layout at Dallas/Ft. Worth International Airport.

The four DFW data sets analyzed are listed in Table 1. These data sets contain over 3500 tracks of which about half are arrivals. The weather for all data sets was VMC with reported visibility's exceeding 10 nm. However for data set DFW.03, a period of IMC weather ended only four hours before the data set began.

Figure 3 depicts the time history of operations for one of the DFW data sets. In this figure, the time of each individual arrival and departure is shown in association with the runway of operation. Several features of the traffic flow can be seen. Note that at approximately 10:10 there is a change in runway configuration - from "north flow" to "south flow". The rate of operations varies greatly with time. Periods of intense activity lasting for 45-60 minutes are followed by lulls in which only modest numbers of operations occur. Such irregularities are attributable primarily to airline scheduling, but they can also be produced by the impact of convective weather upon traffic flow and approach routes.

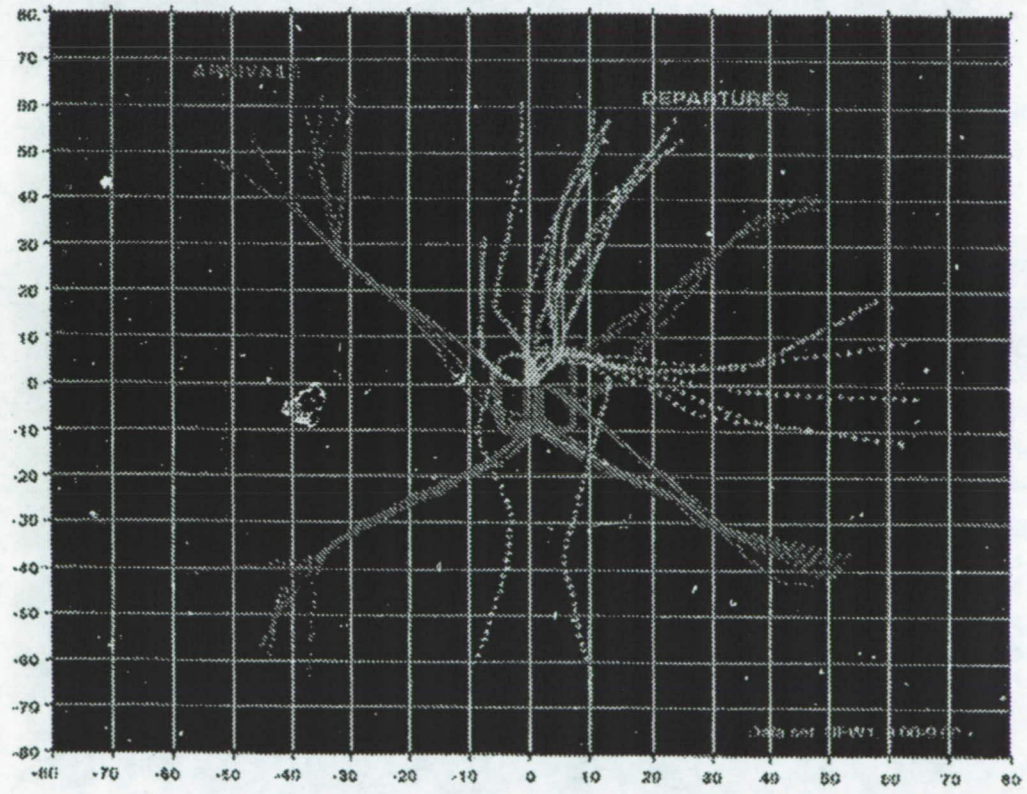

Figure 2. Traffic flow sample at DFW for data set DFW.01.

Table 1. DFW data sets analyzed.

Data Set Name	Date	No. Tracks	Start Time (local 24 hr)
DFW.01	6 DEC 99	826	7:00
DFW.02	10 JAN 00	810	12:35
DFW.03	7 FEB 00	506	17:50
DFW.04	10 FEB 00	1431	8:19

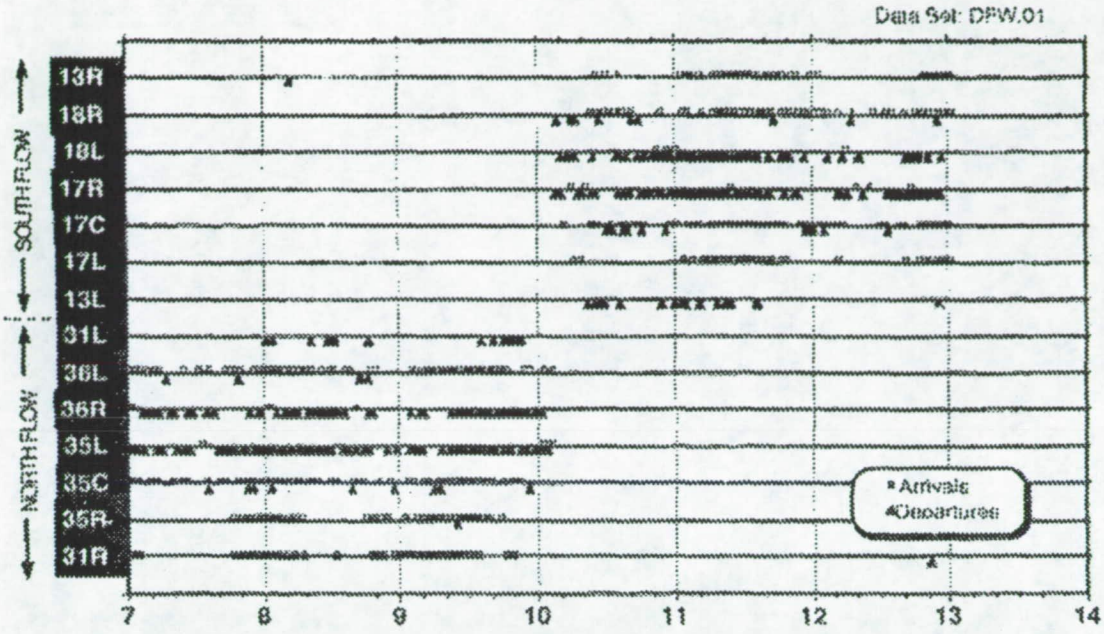

Figure 3. Runway utilization showing a change in runway configuration for data set DFW.01 at 10:30 local time.

ANALYSIS OF LANDING TIME INTERVALS (LTIS)

DEFINITIONS AND BASIC RELATIONSHIPS

While the plots shown in the previous section provide some insight into how the airport was operating, they do not allow us to assess with any quantitative precision the efficiency of the spacings being achieved for any particular runway. In part, this is because the spacing achievable under radar separation standards varies with aircraft weight class, approach speed, approach geometry, and other factors. Techniques for such analyses will now be described. The key feature of the analysis is a focus upon the landing time intervals (LTIs) achieved. The LTI at the runway is defined as the time separation between two successive landings. It is the difference between the time one aircraft crosses the runway threshold and the time the previous landing aircraft crossed the same threshold. (Note: It is also possible to measure LTIs at the outer marker (OM), but such time intervals will not be employed in this paper.).

The throughput of a runway over any arbitrary time period is simply the inverse of the average LTI during that period. For example, if the average LTI is 120 seconds, then the throughput must be 1/120 aircraft/sec or 30 aircraft/hour. We will define the capacity of a runway as the sustained throughput achieved under saturated traffic conditions. Then the mean LTI under saturated conditions is the inverse of the runway capacity.

Figure 4 shows the LTIs for three hours of operations at Dallas/Ft. Worth International Airport (data set DFW.04). Six arrival rushes are clearly seen during the 10 hours of data. This plot provides insight into the extent to which the loading upon the arrival runways was balanced. It can be seen that during the 8:00AM push, runway 35C was not as heavily loaded as the other three runways. But during the 9:30AM push, all four runways appear to have been loaded equally. There were brief periods in which landing intervals of 60 seconds or so were achieved for several successive aircraft.

While we have chosen to measure LTIs at the runway, interarrival times can also be measured at the outer marker. A comparison of the LTI measured at the outer marker and the runway is provided in Figure 5 (using data from data set DFW.01). It can be seen that there does not appear to be any clear tendency for the LTI's to become either greater or smaller between the outer marker and the runway. This implies that control actions taken within the outer marker are not significantly changing the interarrival times.

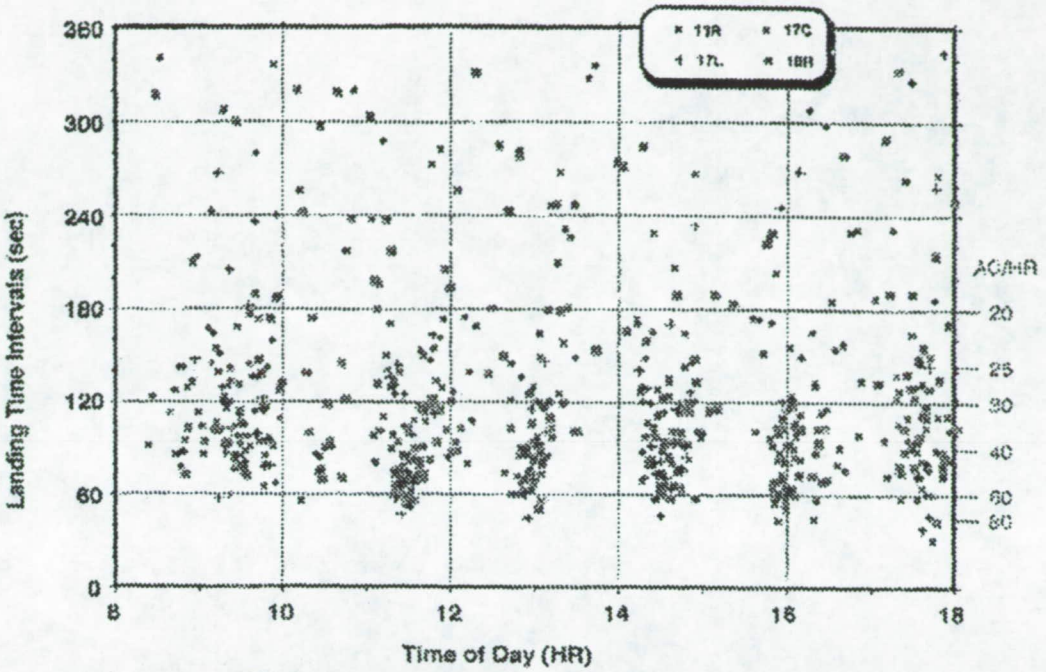

Figure 4. Landing time intervals for data set DFW.04

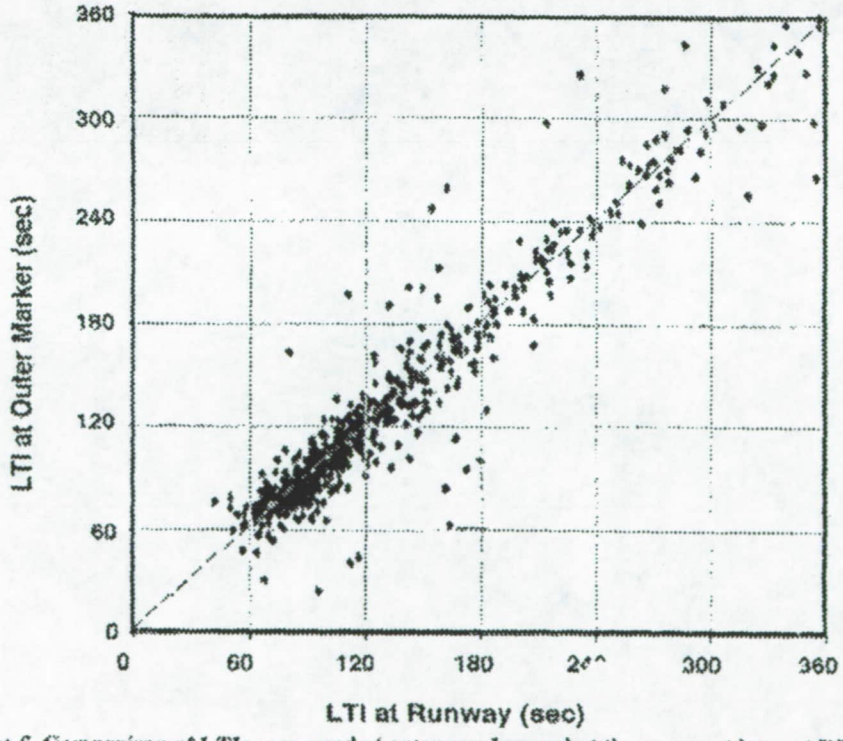

Figure 5. Comparison of LTIs measured at outer marker and at the runway (data set DFW.01).

A histogram of the observed landing time intervals for the four combined data sets is provided in Figure 6. The most common separation was in the 90-100 second range. Figure 7 provides a similar histogram for the minimum in-trail separations observed for 801 arrival pairs in which both aircraft were in the "large" weight class. Separations below 2 nmi appear to be mostly due to the occasional use of visual procedures in which altitude separation was maintained visually. In both figures a line showing a theoretical fit to the histogram is shown. (An explanation of the theory follows.)

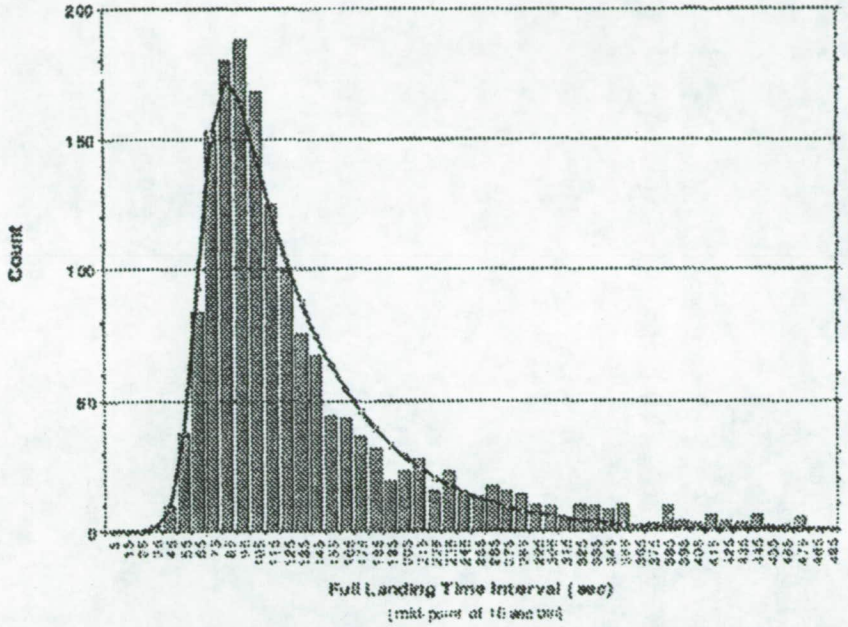

Figure 6. Landing time interval histogram for combined DFW arrivals (1758 pairs).

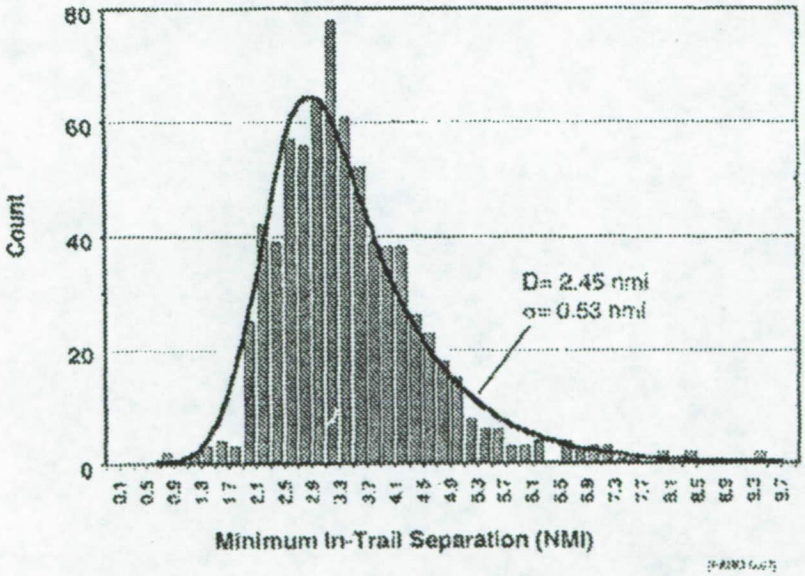

Figure 7. Minimum in-trail separation observed for 801 pairs with both aircraft in the large weight class (DFW).

VANDEVENNE MODEL FOR INTERARRIVAL SEPARATIONS

Any analysis of actual runway operations must recognize that at some times the traffic flow will be less than the runway capacity and that gaps will occur between aircraft that are not due to any inefficiency in the spacing process. A statistical model that takes this into account helps avoid confusing these gaps with excess spacing inserted by the final spacing process. Vandevenne [5] developed such a model for the distribution of observed interarrival separations. PARO employs the Vandevenne model to provide a more robust analysis of final spacing performance. This section describes that model.

The Vandevenne model is motivated as follows: Let us assume that controllers attempt to achieve an interarrival time separation D that represents the closest comfortable target spacing for specified separation standards and operational conditions. The actual time separation achieved, S, will differ from D for two reasons. First, there is imprecision in spacing. Second, there may be gaps in the arrival stream that are too large to be closed by the level of control available. The Vandevenne model assumes that the errors and gaps are additive so that

$$S = D + \varepsilon + g \qquad (1)$$

where ε is the imprecision error and g is the time gap that cannot be closed. The model assumes that ε is normally distributed according to $N[0, \sigma^2]$.

Vandevenne noted that if the arrival stream is random at a given average arrival rate λ, the time gaps between arrivals prior to application of any control actions will have a Poisson distribution such that

$$f_g(x) = \lambda \exp(-\lambda x), \quad x \geq 0 \qquad (2)$$

It should be noted that although time separations in a single arrival stream will not be random because of in-trail separation standards, the merging of multiple independent streams results in an initial set of interarrival times that is approximately Poisson. Vandevenne assumed that all interarrival spacings will include a time gap component, and that this time gap will have a Poisson distribution. The components of S are summarized in Table 2.

Vandevenne showed that the resulting probability density function for S is

$$f_s(y) = \lambda \exp\left[-\lambda\left(s - D - \frac{\lambda\sigma^2}{2}\right)\right] F_{SN}\left(\frac{s - D - \lambda\sigma^2}{\sigma}\right) \qquad (3)$$

where F_{SN} is the standard normal distribution.

In many analyses of actual data, the value of λ changes during the period of observation. This violates the assumptions in the Vandevenne model. For that reason, λ should not be viewed as providing a good indication of the actual arrival flow rate in the data. It is better to view it as merely a parameter of the distribution that is used to correct for the existence of time gaps in the interarrival time observations.

Table 2 The Vandevenne Model

Variable	Definition	Distribution
D	Time separation that controller attempts to achieve.	Fixed for a given aircraft pair
ε	Error in achieving targeted time separation	normal, zero mean $f_\varepsilon(x) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{x^2}{2\sigma^2}\right)$
g	Time gaps in arrival stream that cannot be closed by control in terminal area.	Poisson $f_g(x) = \lambda \exp(-\lambda x)$

The targeted time spacing, D, is a key parameter since the inverse of D is the inherent capacity of the runway. Note that when unsaturated flow exists, the mean observed spacing can be significantly greater than D. A positive bias results if the mean spacing is assumed to be equal to the targeted spacing. The Vandevenne model can produce a nearly unbiased estimate of D under such conditions. This results in a more robust analysis that is better suited to automated processing. Experience has shown that the form of the Vandevenne model provides a good fit to actual data. It has the essential characteristic of a major peak reflecting the predominant normally distributed errors and a long tail reflecting gaps arising from other processes.

INTERARRIVAL SEPARATION AND CAPACITY

We will now discuss how the parameters of the Vandevenne distribution relate to runway capacity and to potential FAST benefits.

The capacity of a runway (defined as the sustainable throughput when saturated with traffic) is approximately 1/D. FAST capacity benefits are assumed to be derived from reductions in the value of D.

At first glance, it appears that the parameter σ has no effect on capacity since the spacing error it produces tends to average to zero. However, it is commonly assumed that in actual operations the value of D is affected by σ because of a need to insert a safety buffer between each pair of aircraft. This buffer guarantees that imprecision will not cause frequent violations of separation standards. The size of the buffer is selected to keep the rate of separation violations below some level, α. If σ is decreased, the safety buffer can be decreased. For the Vandevenne model, we can model the target separation as

$$D = D_0 + \sigma F_{std}^{-1}(1-\alpha) \qquad (4)$$

where D_0 is the required minimum time separation, F_{std}^{-1} is the inverse of the standard normal distribution (with zero mean), and α is the allowed rate of violating this separation. Figure 8 shows how the capacity of a runway is affected by the value of σ when the uncertainty buffer corresponds to either 2σ or 3σ.

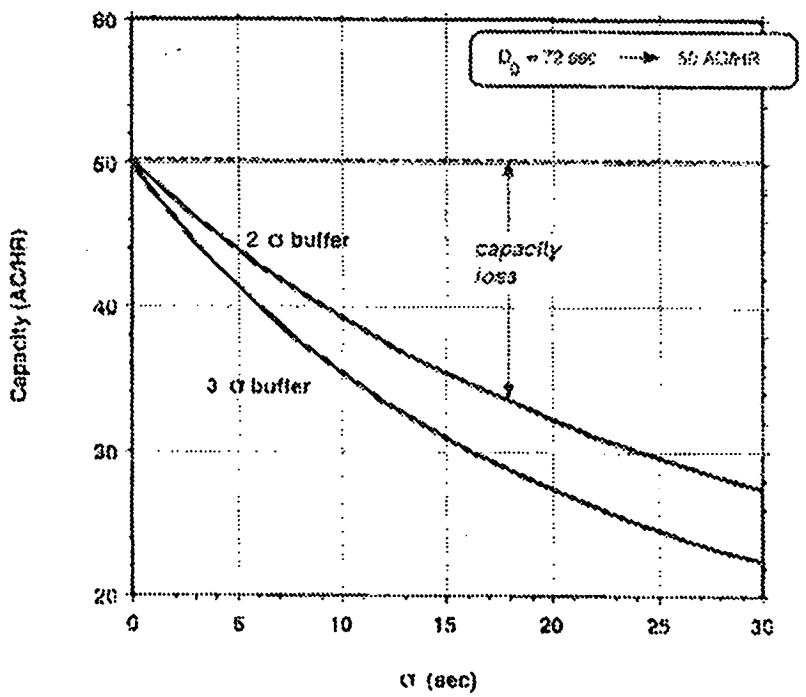

Figure 8. Effect of imprecision (σ) upon runway capacity.

In many cases, the major determinant of D_0 is the in-trail wake vortex separation standard. This standard depends upon the aircraft weight class combination for a pair of successive arrivals. In translating the distance standard to an equivalent time standard, we must also consider the speed profiles of the aircraft on final approach.

To provide a more relevant comparison of aircraft with different weight classes and speeds, we will usually subtract the computed separation standard from the observed separation to yield the excess separation S^* defined as $S^* = S - D_0$.

When the value of S^* is negative it means that the actual separation achieved was less than that indicated by the applicable radar separation standard. This does not necessarily mean that any standards were violated since under visual meteorological conditions the radar separation standards do not have to be applied to aircraft that have their traffic in sight.

The advantage of using S^* is that it allows combining pair separations values for all aircraft types under the assumption that the applicable values of σ and λ are independent of an aircraft's weight class and final approach speed. This assumption appears justified by data analysis completed to date for Dallas/Ft. Worth, but should be verified again when different airports are analyzed.

MAXIMUM LIKELIHOOD ESTIMATION OF MODEL PARAMETERS

Given a set of interarrival time separations, how do we go about finding the model parameters for fitting the Vandevenne model to the data? Vandevenne suggested using a maximum likelihood estimation technique. Suppose that we observe N arrival pairs. Let the i^{th} pair have separation y_i. Then the log likelihood function is

$$L = \sum_{i=1}^{N} \ln[f_y(y_i)] \qquad (5)$$

where f_y is the probability density function for the separation. The maximum likelihood set of parameters is the set that maximizes this function. Vandevenne found the maximum likelihood values by generating contours of likelihood and using search techniques on these contours. While this method is theoretically sound, the estimation of the likelihood function for each point on a contour involves N separate evaluations of the density function f_y. If large databases containing tens of thousands of arrivals are to be analyzed, the computational load could be a hindrance to the analysis. For this reason, an alternative technique was developed that computes an approximate likelihood value directly from the histogram. For this technique, the N data points are compiled into a histogram with H bins. The likelihood factor is calculated for a separation at the midpoint of the histogram bin. The same factor is then assumed to apply to each point in the bin. For example, let the count of separations falling into bin i be n_i. Let the mid-point of the separation interval for bin i be $\bar{y}_i$. Then the approximate likelihood function can be written

$$L = \sum_{i=1}^{H} n_i \ln[f_y(\bar{y}_i)] \qquad (6)$$

With this approach, the number of times the f_y function must be computed is equal to the number of histogram bins instead of the number of points within those bins. This greatly expedites the search for the maximum likelihood values. Inspection of several cases indicates that as long as the histogram bin width is less than approximately one-half σ, the maximum likelihood parameters derived in this way are almost indistinguishable from those derived by using all N original data points.

What is the accuracy with which PARO is able to estimate the three parameters of the Vandevenne distribution? Clearly, the accuracy will depend on the number of points that are available for forming the estimate. It will also depend upon the bin size used in the histogram. Table 3 presents simulation results for a case in which the true parameter values are D = 72.0 seconds, σ = 18.0 seconds, and λ = 90/HR. For each entry, Monte Carlo simulation was used to generate 100 histograms, each with a bin size of 10 seconds. The standard deviations of the parameter estimation error decreases roughly as the inverse of the square root of the number of points used to construct the histogram - an expected result. For σ, a standard deviation of error that is less than 10 percent of the true value is achieved with 400 data points.

Analysis of LTIs from DFW

The LTI distributions that exist in the four DFW data sets were analyzed by first generating LTI histograms for each set separately and then for the combined data. The maximum likelihood fit of the Vandevenne distribution to each histogram was computed. Figure 9 depicts the histogram of excess LTIs for all four data sets combined.

Table 3. Landing Time Interval (LTI) Analysis

No. Points in Histogram	Mean Error in D	Std. Dev.of D	Mean Error in σ	Std. Dev. of σ	Mean Error in λ	Std dev of λ
200	-0.254	3.249	-0.451	2.730	1.954	10.239
400	0.307	2.483	0.107	1.781	2.562	7.150
800	-0.256	1.774	-0.216	1.300	0.471	4.723
1600	-0.075	1.009	-0.008	0.810	0.536	3.303

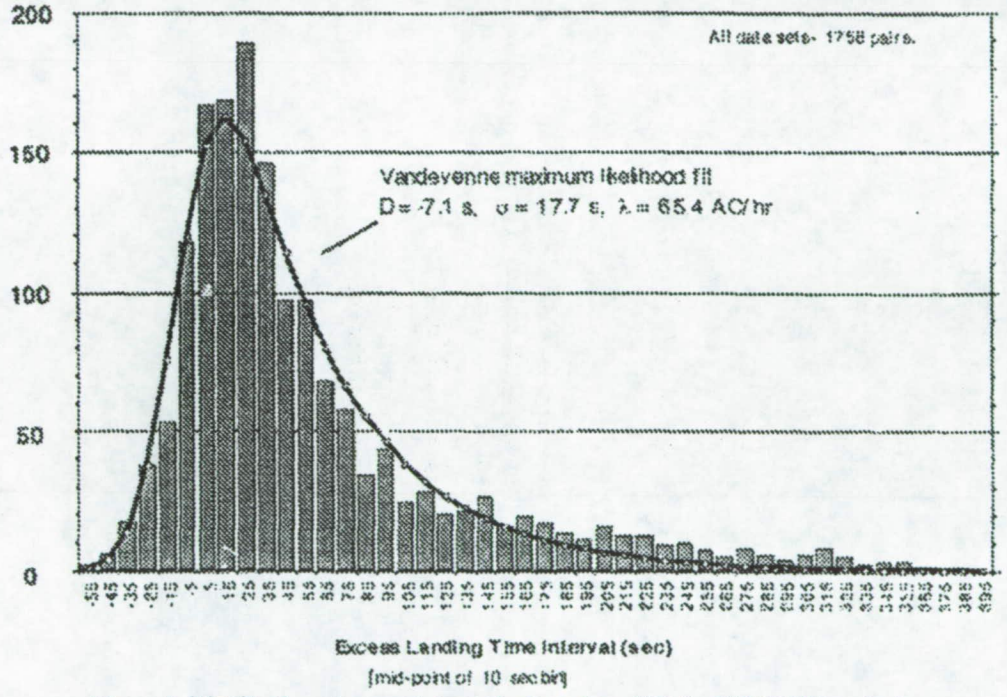

Figure 9. Distribution of excess interarrival spacings (S*) for 1758 DFW arrivals.

The results of the analysis are summarized in Table 4. It can be seen from the histograms that the shape of the distribution closely resembles the Vandevenne distribution. The maximum likelihood parameters are similar for all sets. The combined value of σ was 17.7 seconds. This is only slightly less than the 19-20 second values reported by Ballin and Erzberger. [2]

The fact that D tends to be slightly less than zero means that the aircraft were often achieving intervals smaller than would be possible under radar separation standards. This implies for this predominantly VMC data, it is unlikely that an AFAST calibrated to preserve radar separation standards could have increased throughput by reducing the "imprecision buffers" incorporated in D. Nevertheless, AFAST might have been able to provide benefits by anticipating and removing the larger interarrival gaps that are related to the overall flow to the runways. Additional data analysis is being pursued to confirm this and to address the same question for IMC conditions.

ANALYSIS OF FACTORS AFFECTING SEPARATIONS

This section presents the results of several types of analysis conducted to investigate the reasons for the differences between the LTIs obtained under different conditions.

Correlation Analysis for LTIs

One way of searching for factors that affect final spacing efficiency is to examine the linear correlation coefficient between various variables and the excess separation. This type of analysis can fail to detect certain types of nonlinear dependencies, but will nevertheless identify a number of relationships that deserve further scrutiny.

Figure 10 shows the parameters used to characterize the final approach geometry. The parameters are defined in the runway coordinate system for which the origin is the runway threshold. The parameter YBASE is the y value at which the base leg was established. The parameter yCL2 is the y value at which the aircraft achieved flight along the centerline of the runway. The criteria for "centerline" status is that the aircraft has to be within 2 nmi of the centerline and have a heading within 15 degrees of the runway heading.

Table 4. Errors in Estimation of Vandevenne Parameters

Data Set	No. of Aircraft Pairs	D (sec)	σ (sec)	λ (AC/hr)
DFW.01	418	-7.9	14.3	62.8
DFW.02	391	-0.3	17.2	69.0
DFW.03	242	-0.7	19.9	68.1
DFW.04	707	-12.3	17.6	64.4
All sets combined	1758	-7.1	17.7	65.4

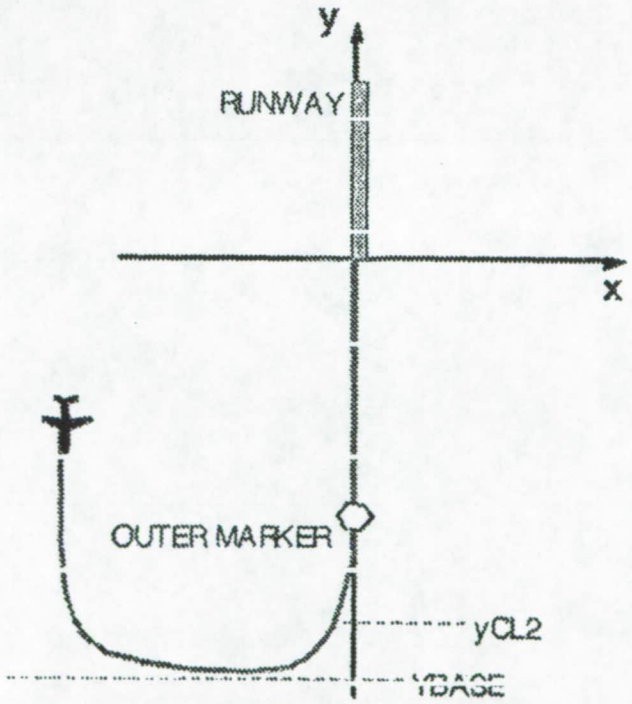

Figure 10. Definition of approach geometry using YBASE and yCL2.

Table 5 defines the "W variables" that were used to describe the pair of aircraft generating a single LTI value.

Table 6 provides the serial correlation coefficient ρ, for these variables when correlated against the full value of the LTI. Only large/large weight class pairs were used to avoid variations due to differing wake vortex separation requirements. If the initial gap between arrivals was too large to be closed by typical control actions, the pair was excluded from this analysis. In this table, p is the probability that the observed correlation coefficient would be as far from zero as observed if the actual value were in fact zero.

Table 5. Descriptive "W Variables" for an Arrival Pair

Variable, W	Definition
absolute LTI	Absolute value of landing time interval (sec).
AP_PATTERN	Approach pattern type of follower (1000=straight-in, 2000=downwind/base)
AZ_FIRST	Azimuth at which aircraft first appeared (degrees).
LTI_OM	Landing time interval at outer marker (sec)
LTImin	Minimum LTI permitted by separation standards (sec).
pathlength	Total path flown in terminal airspace by follower (meters)
S*	Excess landing time interval at runway (sec)
t_on_CL	Time spent "on centerline" (CL) state before landing (sec)
V2/V1	Speed ratio (final) of follower to lead aircraft.
vel_op2	Final speed (at landing) of follower (KT).
vOM2	Velocity of follower at outer marker (KT).
WEIGHT_CL_DIF	Weight class code of lead minus that of follower,
wtclass_lead	Weight class code of lead (1=light, 2=large, 3=B757, 4=heavy)
YBASE	y coordinate of base segment for follower (meters).
yCL2	Centerline intercept coordinate of follower (meters)
yCL2-yCL1	Difference in centerline intercept coordinate of follower and lead (meters).

Table 6. Linear Correlation Analysis of Full LTI: S vs. Variable W, Large/Large Weight Class Pairs

Variable, W	n	mean S	Std. Deviation of S	mean W	Std. Deviation of W	ρ	p	signi-ficance
V2/V1	1211	105.86	42.890	1.014	0.148	-0.149	0.00000	•••
vOM2	1211	105.86	42.890	86.3	14.2	-0.120	0.00003	•••
vel_op2	1211	105.86	42.890	67.221	6.287	-0.082	0.00459	•••
AZ_FIRST2	1211	105.86	42.890	299.6	455.8	-0.045	0.11802	
yCL2-yCL1	1211	105.86	42.890	245.8	9371.4	0.011	0.70289	
AP_PATTERN	1211	105.86	42.890	1067.5	263.3	0.020	0.49171	
WEIGHT_CL_DIF	1211	105.86	42.890	0.015	0.768	0.050	0.08173	•
wtclass1	1211	105.86	42.890	2.045	0.609	0.092	0.00143	•••
LTImin	1211	105.86	42.890	0.4	18.0	0.131	0.00001	•••
yCL2	1211	105.86	42.890	-17056.9	9231.8	0.136	0.00000	•••
pathlength2	1211	105.86	42.890	150929.8	23251.4	0.149	0.00000	•••
YBASE2	1043	107.42	44.733	-17325.3	5428.3	0.157	0.00000	•••
S*	1211	105.86	42.890	29.0	43.6	0.905	0.00000	•••
LTI_OM	1211	105.86	42.890	105.9	45.0	0.911	0.00000	•••

Significance code: • = significant at 0.10 level, •• = at 0.05 level, ••• at 0.01 level

The following observations apply to Table 6:

- Separation increases when the total path length flown increases (ρ = 0.147). This may reflect greater constraints and traffic interactions encountered by aircraft that fly longer paths. It could also reflect the fact that having to maneuver within terminal airspace introduces imprecision into the spacing.

- There is high correlation (ρ = 0.880) between LTI measured at the runway and LTI measured at outer marker. This suggests that if efficient spacing isn't achieved at the outer marker, then it is unlikely to improve much due to actions taken within the marker.

Table 7 provides a correlation analysis for W variables correlated against the excess landing time interval, S*. Here all weight classes can be combined. Note that

- Excess separation is negatively correlated (ρ = -0.203) with weight class difference (lead minus follower). This indicates that when the lead aircraft is heavier, the separation relative to wake separation standards is less.

- Excess separation increases when centerline intercept is closer to the runway (ρ=0.135 for yCL2). Excess separation decreases with more time spent on the centerline (ρ=-0.122 for t_on_CL). The reason for this is not clear, but may have something to do with the ability to tighten separation through speed control as compared to trying to achieve tight separation by a precise turn from a short base leg.

- Excess separation increases when pathlength increases. (See earlier comment for Table 6).

- There is negative correlation (ρ = -0.254) with absolute LTI allowed by separation standards. This suggests that there is a tendency to space closer than the standard for the larger standards, perhaps through use of VMC procedures.

There is high correlation (ρ = 0.905) with excess LTI measured at outer marker. Again, this indicates that actions taken after the outer marker have little impact on the final time separations.

ADDITIONAL ANALYSIS OF FACTORS AFFECTING LTIS

Differences Between Runways

An obvious question to ask is whether LTI distributions are the same for all runways. Figure 11 provides histograms of LTIs for each runway at DFW using the combined data sets. The LTI distribution for each runway appears to be similar except possibly for runway 18L for which very few arrivals were observed.

Table 7. Linear Correlation Analysis of Excess LTI: S* vs. Variable W, all aircraft pairs

Variable, W	n	mean S	Std. Deviation of S	mean W	Std. Deviation of W	ρ	p	signi-ficance
V2/V1	787	101.36	36.349	1.017	0.154	-0.138	0.00011	•••
vel_op2	787	101.36	36.349	67.672	6.245	-0.123	0.00056	•••
vOM2	787	101.36	36.349	86.7	13.6	-0.035	0.32483	
AZ_FIRST2	787	101.36	36.349	284.5	448.8	-0.025	0.48288	
yCL2-yCL1	787	101.36	36.349	167.1	9450.6	0.035	0.32807	
AP_PATTERN	787	101.36	36.349	1079.9	281.5	0.048	0.17857	
LTImin	787	101.36	36.349	-0.2	17.6	0.116	0.00121	•••
yCL2	787	101.36	36.349	-16967.1	9141.7	0.144	0.00006	•••
pathlength2	787	101.36	36.349	152726.2	22616.6	0.147	0.00004	•••
YBASE2	688	102.98	37.131	-17549.5	5447.4	0.179	0.00000	•••
LTI_OM	787	101.36	36.349	101.7	37.9	0.880	0.00000	•••
S*	787	101.36	35.349	31.8	35.2	0.955	0.00000	•••

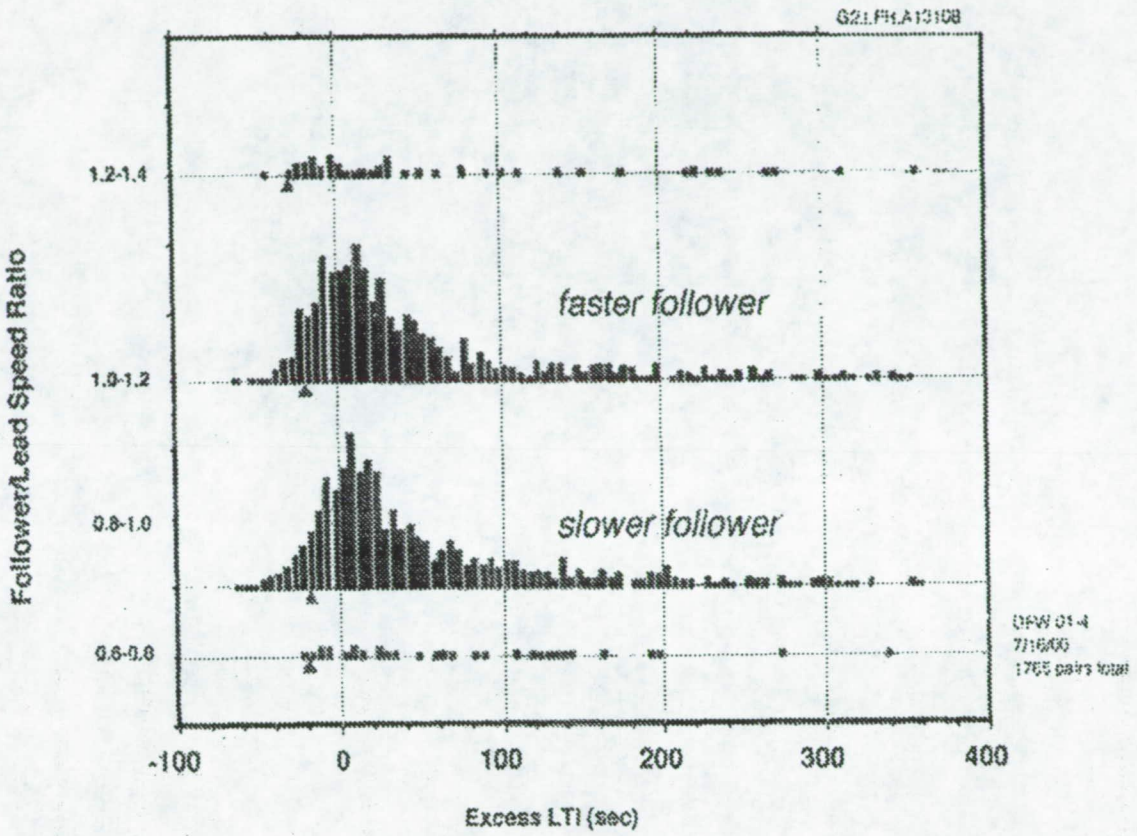

Figure 11. LTI differences between runways

Figure 12 examines the effect on the LTI distribution of the speed ratio between the lead and following aircraft. If it is more difficult for controllers to anticipate the effect of such speed differences, then differences might be expected to appear in the histograms. For the data in Figure 12, the values of the targeted time separation, D, are quite similar except possibly for lower D value for the case of a following aircraft more than 20% faster than the lead aircraft. In general, it appears that controllers at DFW are quite skilled at taking the differing landings speeds of aircraft into account when spacing them.

Approach Patterns

It seems possible that the precision of interarrival spacing can be affected by the geometry available for making the final spacing adjustment. For aircraft that fly a downwind segment, the controller is able to choose the location of the base segment to achieve proper spacing. But the turn required doing so can be a source of imprecision. Is there a difference in spacing performance that should be addressed by tools such as FAST?

An analysis of approach patterns was conducted by dividing all pairs of successive aircraft into four groups depending upon whether or not the approach involved a downwind-base trajectory. Approaches without a downwind phase were called "straight-in", although it should be noted that some of these aircraft approached the runway centerline at a fairly large angle. The set labeled "downwind->straight" includes all pairs for which an aircraft on straight-in approach was followed by an aircraft flying a downwind segment. The four histograms that result are shown in Figure 13. At DFW, aircraft are generally first directed to the cornerpost fixes that are most consistent with the north/south direction of flow, and hence straight-in patterns predominate for final approach. While there are no statistically significant differences in the S^2 distribution for this data, it does appear that aircraft following a downwind leader tend to have higher median values

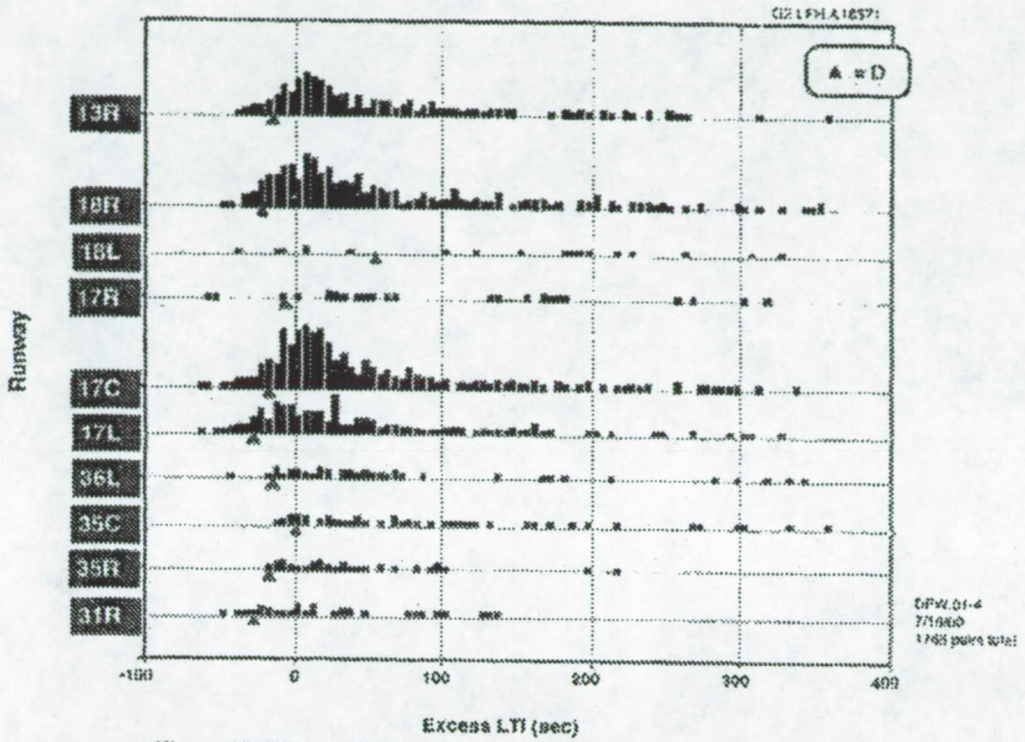

Figure 12. Effect of follower/leader speed ratio on landing time intervals

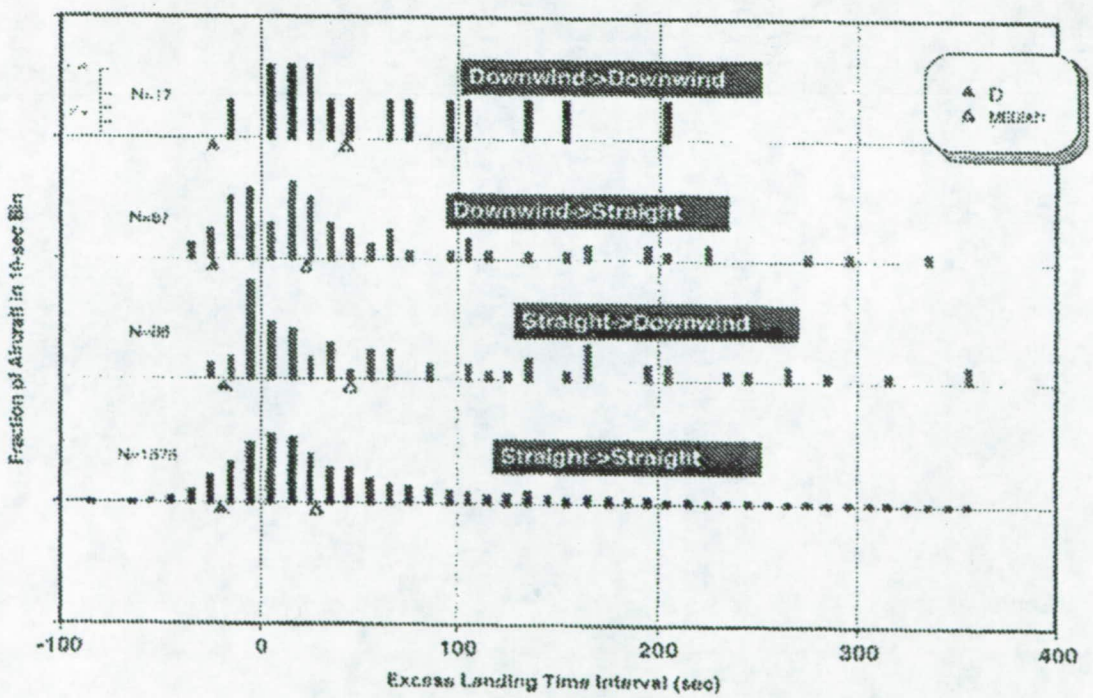

Figure 13. Excess landing time for different approach pair types

One additional question concerns the relationship between time intervals and actual in-trail spacing. In this report, we have expressed separations in terms of time, but actual separation standards are expressed in distance. Figure 14 shows the relationship between excess landing time intervals and excess in-trail spacings for arrival pairs in data set DFW.01. It can be seen that there is high correlation (ρ = 0.843) between the two measures. The dotted line shows the space-to-time conversion that would if the excess spacing is traversed by the trailing aircraft at a speed of 150 knots. This analysis suggests that conclusions about performance deduced from inspection of time separations are likely to be the same as those of an analysis that used only spatial separations.

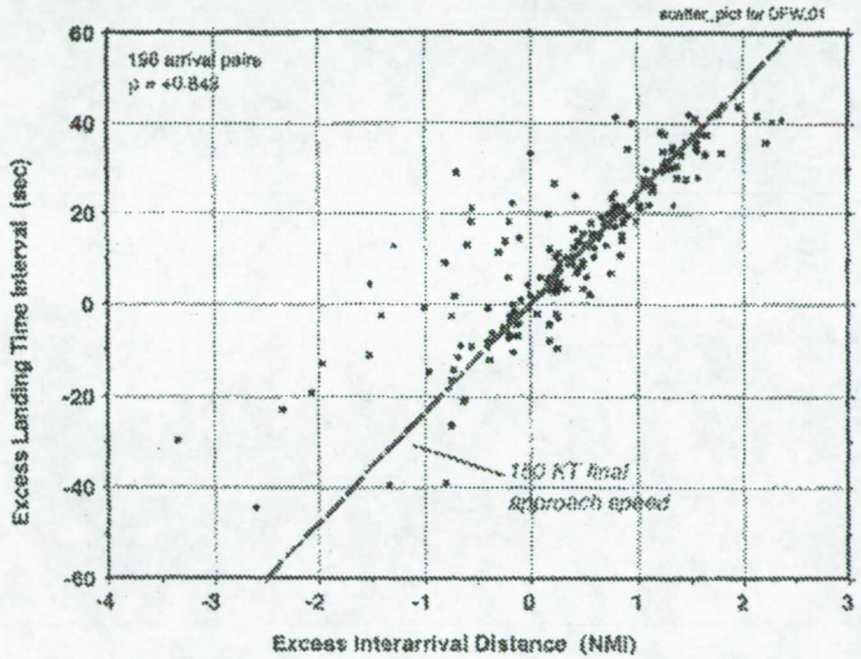

Figure 14. Comparison of excess spacing expressed in time and distance

CONCLUSIONS AND RECOMMENDATIONS

This work has developed robust, computationally practical data analysis routines that can be used to provide insights into runway operations through analysis of radar data. The questions that can be addressed are relevant to the benefits mechanisms of AFAST and to determining the total benefits that can be achieved with AFAST implementation.

The following observations apply to the four data sets analyzed for DFW. It should be noted that all these data sets involved VMC weather and the generality of the conclusions drawn from this limited data set has not been proven.

- In short rush periods (of 10 minutes or so), very high landing rates of 60 AC/HR or more are obtained on single runways. It is not clear that the high peak rates observed can be obtained in IMC. Nor is it clear that they can be sustained in periods of prolonged saturation.
- In general, the LTI achieved at the outer marker is preserved at the runway. Variations appear to be random with no discernible tendency to change in a given manner. Thus, there is little evidence that visual separation practices applied within the outer marker are having a significant impact upon spacings.
- The targeted separation, D, appears to be about 72 seconds, which corresponds to a throughput of about 50 aircraft/hour. The fact that this rate is almost never sustained in practice suggests that there is an opportunity to increase throughput if the consistency of flow to the final vector position is improved.

- Targeted separation (D) tends to be slightly less (by about 7 seconds or 0.3 nmi) than the value that would be expected if radar separation standards were the sole determinant of target separation. This suggests that visual procedures near the runway may have allowed separations to be tightened enough to overcome the effects of any "safety buffers" that were applied during the earlier radar separation process.

- The occasional presence of larger LTIs during periods of saturated flow is a further indication that irregularities in the flow may be contributing to a loss of throughput in VMC. This phenomenon deserves further study since the ability of AFAST to reduce such irregularities provide capacity benefits under VMC conditions.

REFERENCES

1. Boswell, Steven B., "Evaluation of the Capacity and Delay Benefits of Terminal Air Traffic Control Automation," Project Report ATC-192, Report No. DOT/FAA/RD-92/28, MIT Lincoln Laboratory, 14 April 1993.

2. Ballin, Mark G., and Erzberger, Heinz, "Benefits Analysis of Terminal-Area Air Traffic Automation at the Dallas/Fort Worth International Airport," AIAA 96-3723, AIAA Guidance, Navigation and Control conference, San Diego, July 29-31, 1996.

3. Weidner, T., Couluris, G.J., and Hunter, C.G., "CTAS Benefits Extrapolation Preliminary Analysis," 98156-01, Seagull Technology Inc., February 1998.

4. Thompson, S.D., "Terminal Area separation Standards: Historical Development, Current Standards, and Processes for Change," Project Report ATC-258, MIT Lincoln Laboratory, 16 January 1997.

5. Vandevenne, H.F., and Lippert, M.A., "Using Maximum Likelihood Estimation to Determine Statistical Model Parameters for Landing Time Separations", 92PM-AATT-0006, 17 March 2000 [Originally published as internal memorandum 27 April 1992].

ASSESSING DELAY BENEFITS OF THE FINAL APPROACH SPACING TOOL (FAST)*

Jerry D. Welch and John W. Andrews
MIT Lincoln Laboratory
244 Wood Street
Lexington, MA 02420-9108

John E. Robinson III
NASA Ames Research Center
Moffett Field, CA 94035-1000

ABSTRACT

Air traffic delays are costly. NASA is developing the Final Approach Spacing Tool (FAST) to help increase throughput and reduce the approach component of airborne delay. Analysis and field trials have suggested that FAST can help controllers increase arrival throughput on busy runways by several aircraft per hour.

Two major simulation studies have predicted that delay reductions from such throughput increases would save several hundred million dollars annually. The studies also provided useful data on operations at major airports. However, their predictions disagree on delay savings for some airports and omit other airports of interest. Their predicted delay savings for some airports are higher than actual reported delays for those airports. Because of resource and time limitations neither study considered storm disruptions to arrival routes, and neither addressed downstream delay propagation caused by schedule disruption. Both of these effects change the dollar savings from FAST. Although delay propagation can only multiply delay savings, the effect of a storm can be positive or negative. Route disruptions from storms can temporarily prevent FAST from operating. But storms can also leave large queues that magnify the value of incremental capacity increases from FAST when it returns to operation.

In this paper we summarize and compare the two benefit models. We present a simple benefit model for judging the accuracy of the two models and for helping to ranking benefits among airports. We examine measured delay data published after the completion of the studies that helps to validate the model results. We use this data to examine delay coupling between airports and we use it with a published model for delay propagation to estimate downstream delay multiplication for Dallas Fort Worth Airport. Finally, we examine how storms affect delay and how they might modify the estimated FAST benefits.

*This work was performed for the National Aeronautics and Space Administration under Air Force Contract No. F19628-00-C-0002.

INTRODUCTION

The cost of air traffic delay grows each year as delays and demand increase and as fuel costs increase. The Air Transport Association estimated that airborne delays cost US air carriers about $800M in 1999.[1]

Figure 1 illustrates airborne delay during congested periods at DFW on Monday10 April 2000. The figure shows the arrival rate at 15-minute intervals and the corresponding average airborne delay for arrivals in each 15-minute period. The arrival rate was determined from radar data, and the average airborne delay was obtained from the FAA's Consolidated Operations and Delay Analysis System (CODAS).[2] Six distinct arrival rushes are evident, and associated with each rush is a transient build-up of queuing delay.

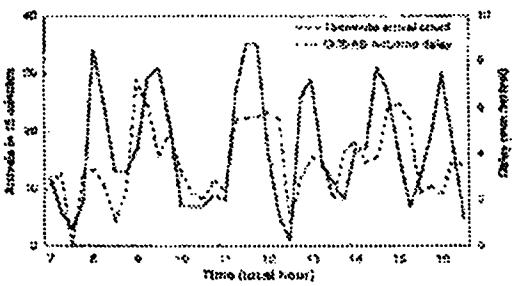

Figure 1 Arrival Rate and Airborne Delay at Dallas on 10 April 2000.

As part of the Center-TRACON Automation System (CTAS), NASA is developing the Final Approach Spacing Tool (FAST) to help reduce this kind of delay.[3] FAST is intended to increase terminal throughput and reduce the approach component of airborne delay caused by queuing congestion at the airport. It accomplishes this by providing planning advisories for efficient runway balancing and arrival sequencing, and by helping to increase the accuracy of final approach spacing.

Analyses, simulations, and field trials indicate that FAST could help controllers increase arrival throughput on a busy runway by several aircraft per hour. Major independent studies by Seagull Technology, Inc. and Logistics Management Institute (LMI) have estimated the potential dollar savings from such throughput improvements at 10 major US airports.[4,5,6,7,8,9] Both studies determined from initial analysis that FAST has the potential to decrease aircraft inter-arrival times by

about 15 seconds and thereby increase throughput by about 4 arrivals per hour per runway. This is accomplished by helping controllers reduce the variance of inter-arrival timing. The two studies assumed that this more precise arrival timing allows controllers to reduce inter-arrival spacing without increasing the incidence of separation violations. They then analyzed the extent to which such a throughput increase would reduce approach delay from congestion during arrival rushes.

The studies employed independent demand and capacity estimates and used separate queuing engines to calculate the reduction in queuing delay. The modeling process is illustrated in Figure 2. A queueing engine converts time-varying airport capacity and demand to delay. Seagull Technology used an FAA-provided discrete event simulation to convert capacity and demand profiles to delay. LMI integrated the Kolmogorov queuing equations to derive delay profiles. They used independent, but similar, models of direct operating cost to aggregate and convert aircraft arrival delay savings to dollars.

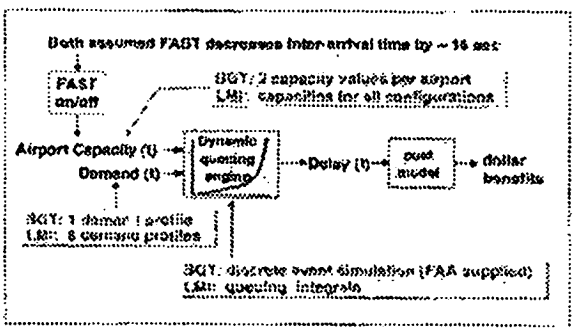

Figure 2. Comparison of the Seagull Technology (SGT) and Logistics Management Institute (LMI) Benefit Models.

The models run capacity/demand scenarios for each airport with today's runway capacity, and then repeat each scenario with the increased capacity that would result from reduced inter-arrival times achieved with the help of FAST. Seagull Technology modeled only 2 capacities for each airport, whereas the LMI capacity model accounted for all major runway configurations at each airport. The Seagull Technology demand model used a single demand profile for each airport, whereas the LMI model used 6 demand profiles per airport to account for seasonal differences and differences between weekday and weekend schedules. The Seagull Technology study examined 35 airports. The LMI study analyzed 10 airports.

Both studies adopted similar limitations of scope. They did not validate their predictions with actual delay measurements for the airports studied. Their models focused on delays in IMC and VMC and did not reflect the fact that FAST cannot predict flight trajectories and must cease operating when hazardous weather disrupts arrival routes. It follows that they also did not account for the large residual queues that must be serviced when routes re-open after storms.

The dollar savings estimates of both were kept intentionally conservative by focusing only on direct operating cost savings. Downstream delay propagation costs, which can be large for days with unusually adverse weather or high demand, were omitted for lack of a cost model applicable to all airports.

The ten airports modeled by LMI were all included in the Seagull Technology study. The overall cost savings predicted for the ten airports by Seagull Technology was $375M/year. LMI estimated a savings of $319M/year. The dollar estimates for some airports differ more significantly than do these overall numbers. Figure 3 compares the results for the 10 airports that were studied in common.

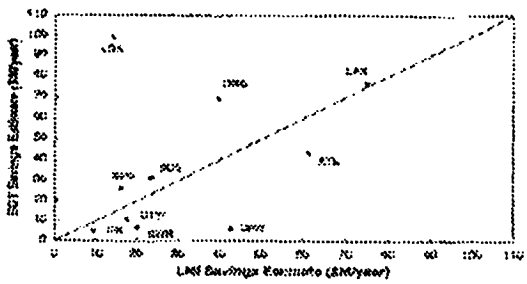

Figure 3. Comparison of Seagull Technology and Logistics Management Institute Benefit Estimates.

Ideally, the points for all of the airports would fall near the line of unit slope included in the figure. The disagreement is large for Chicago, and Atlanta. It is very large for LaGuardia and Dallas/Fort Worth. These differences make it difficult to determine relative benefit rankings for the airports.

A SIMPLE RANKING RULE

To help resolve these differences we developed a simple benefit model based on steady state queuing. These benefit analyses reduce ultimately to calculations of the change in average annual delay that results from a given capacity increase at all runways. Although steady state queueing theory does not allow exact delay calculations when capacity and demand vary with time, it does provide insight into average annual delay trends.

If we assume that, on average, all airport runways are loaded with the same ratio of demand to capacity, and if the capacity of all runways is increased identically, then

to first order, steady-state queuing tells us that the delay reduction at each airport will be proportional to N^2/R, where N is the traffic count and R is the total number of runways at the airport.

This prediction was tested against the results of both studies. Figure 4 shows the LMI annual savings estimate for each of the 10 airports plotted as a function of N^2/R. The data label for each airport includes in parentheses the number of runways capable of handling commercial flights at that airport.

The LMI benefit numbers follow the N^2/R trend reasonably well. Also plotted is a linear least-squares fit to the savings for the 10 airports, which can be used to estimate savings for other airports based on their operations and runway counts. We estimated the savings for the FAA Free Flight Phase-1 airports that were not included in the LMI study (Philadelphia, Charlotte, Denver, Miami, Minneapolis/St. Paul, and St. Louis). If the LMI savings estimates for FAST at the 10 study airports are correct, the benefits for the six additional airports range from $11M per year for Denver to $35M per year for Miami.

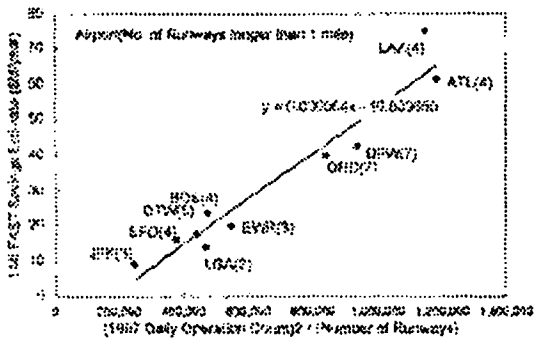

Figure 4. LMI FAST savings estimates vs. (N^2/R).

Overall, the results indicate that airlines using the four largest airports (LAX, ATL, DFW, and ORD) would save the most from FAST. If controllers using FAST could indeed increase runway arrival throughput by 4 aircraft per hour at all 10 airports, airline operators would recover an estimated $460M annually in direct operating costs.

CODAS DELAY DATA

We now examine the use of recently available delay measurements to validate the model results. When the two benefits studies began, there was no available delay metric commensurate with the queuing delay considered in the studies. For a number of years the main quantitative delay reports have been those based on FAA OPSNET data, which counts flights with schedule delays exceeding 15 minutes.[19] In 1977 the FAA began to make available statistics on delays of all magnitudes as part of the Consolidated Operations and Delay Analysis System (CODAS).[2]

CODAS delay is gathered and archived for more than 100 airports and provided on the internet for authorized users. The CODAS database includes delays from several phases of flight. Its "arrival" delay and "airborne" delay estimates appear to be most relevant to the estimation of FAST benefits. Both types of delay are averaged every fifteen minutes and reported in units of minutes per arrival. In deriving these averages, the CODAS processing algorithms count early arrivals as having zero delay rather than negative delay. Figure 5 illustrates the differences between these delay metrics.

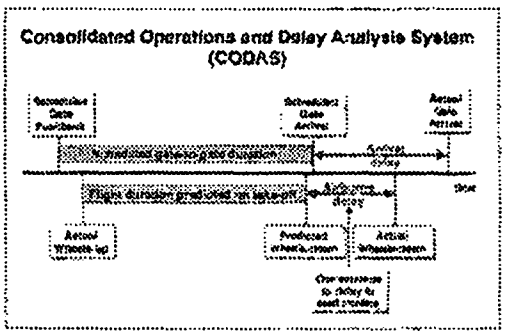

Figure 5 Definition of CODAS Airborne and Arrival delays.

Airborne delay is measured relative to the flight duration predicted at the time of departure. It is the actual flight duration (wheels-up to wheels-down) minus the predicted flight duration. Airborne delay does not include departure delays. The direct operating cost of airborne delay can be readily calculated. Although some airborne delay can be caused by en route weather and traffic flow problems, normally one of its largest components is the terminal queuing delay that runway capacity improvements from FAST are intended to reduce.

Arrival delay is measured relative to scheduled arrival time. It is the actual gate arrival time minus the final airline Computer Reservations System scheduled gate arrival time for the flight. If the flight duration predicted on take-off is the same as the scheduled flight duration, the arrival delay is the sum of the departure delay, the airborne delay, and the taxi-in delay. Arrival delay is usually larger than airborne delay.

Arrival delay is not related to operating cost in a simple way. It includes the delay that results when the traffic flow management process holds aircraft on the ground to minimize airborne delay. But reducing airborne delay by holding aircraft on the ground does not change the fact that aircraft land behind schedule. These schedule delays can result in downstream ripple costs

that are more difficult to account for than the direct operating costs associated with airborne delays. An examination of arrival delay can help determine the full benefit to be expected from an improvement in airport capacity.

MODEL RESULTS AND CODAS DATA

Figure 6 compares the 1997 CODAS average annual airborne delay at ten airports with the LMI model estimates for annual delay that would have been saved in 1997 by FAST and TMA at those airports. The figure shows that the general trends for the LMI savings estimates and the reported delays are similar. However, the LMI estimates for three of the airports would appear to be illogically large in that the delay savings estimates from a small incremental improvement in runway capacity exceed the total reported airborne delay at those airports.

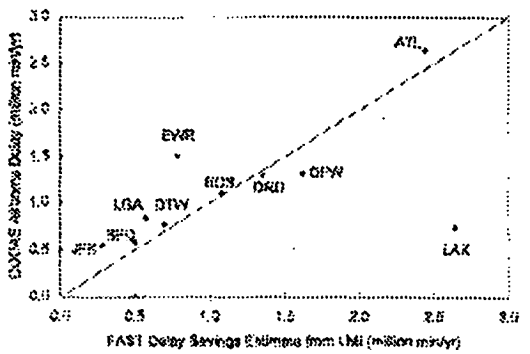

Figure 6. CODAS airborne delay in 1997 vs. LMI estimate of delay savings from FAST.

The Seagull Technology study showed a similar trend. It estimated that five airports would experience delay reductions from FAST that exceeded their CODAS airborne delays, with savings estimates for LAX, LGA, and ORD exceeding CODAS delays by large factors.

Systematic underestimation of airborne delay by CODAS can make delay savings estimates appear high for all airports. Airborne delays are computed relative to operator estimates of flight time at takeoff. When aircraft operators predict flight durations based on mean historical flight times rather than the shortest feasible flight times, CODAS underestimates airborne delay.

This discrepancy is most significant for Los Angeles. Both the Seagull Technology and LMI studies predicted a savings of over 2.5 million minutes/year at LAX, whereas the reported airborne delay at LAX was less than 1 million minutes per year. This is likely caused by traffic flow management procedures. When the air traffic management system uses ground holds and taxi-out delays to regulate and meter the flow of aircraft to an airport, airborne delay is lower than the delay calculated by queuing models assuming random arrivals, and there is a net increase in arrival delay.

CODAS statistics for LAX consistently show a relatively high ratio of arrival delay relative to airborne delay. In 1997, the average arrival delay at LAX was 5.9 times larger than the average airborne delay. The average arrival/airborne delay ratio for the 30 busiest airports was 3.4, and only two of those 30 airports (Phoenix and Washington Dulles) had higher ratios than LAX. At the other extreme, LGA and EWR, which lie farthest above the line of unit slope in Figure 6, had lower than average ratios: 2.1 and 2.4 respectively.

It appears that airborne delay at LAX may be less than at other major airports because the air traffic management system is both motivated and able to regulate and meter en route flow to LAX with unusual consistency. LAX is unique among major US airports in the complexity of its airspace and the uniformity of its weather. Its dense traffic and constrained approach routes provide the motivation for consistent metering to avoid the need for airborne holding. Its freedom from storms makes it possible to meter with consistency. Smoothing of the en route arrival flow further reduces unpredictability from queuing contention at runways. Schedules for flights to LAX can absorb repeatable delays caused by metering and thereby trade longer average flight times for increased schedule predictability.

The high schedule predictability at LAX is evident in the CODAS statistics. Table 1 summarizes arrival delay statistics for seven airports in 1997.

Table 1. CODAS annual arrival delay (minutes/arrival) for seven airports in 1997.

	ATL	BOS	DFW	EWR	LAX	LGA	PHL
Mean	14.5	13.9	10.0	15.9	11.2	10.0	11.0
Std Dev	8.6	11.9	7.9	12.4	5.7	7.2	6.5

In 1997 the standard deviation of CODAS arrival delay at LAX was only 51% of the mean. The percentages for BOS, DFW, EWR, and LGA were respectively 86%, 79%, 78%, and 72%. Increasing schedule predictability at LAX reduces airline cost by decreasing the propagation of schedule delays to downstream airports. This allows more efficient use of crews, ground facilities, and aircraft. We discuss delay propagation more fully below

CODAS delay statistics provide vauable insight into operational differences among airports. They are also useful for assessing changes in delay, correlating delay between airports, and studying the effects of weather on delay, as shown in the following sections.

IMC DELAY

FAST is intended to help improve airport capacity. Transitions from VMC to IMC cause measurable statistical changes in airport capacity. Therefore, quantifying the relationship between local meteorological conditions and measured delay (i.e., using the transition from IMC to VMC as an analytical surrogate for a capacity increase) can provide baseline comparisons for capacity modeling results.

The FAA's CODAS delay database includes local ceiling, visibility, and wind as well as a meteorological condition indicator that switches from IMC to VMC when visual approaches are allowed at each airport. We used this database to examine the dependence of actual delay data on local meteorological conditions at key airports.

CODAS defines Visual Meteorological Conditions as the combination of ceiling and visibility for which visual approaches are allowed. To support visual approaches the ceiling must be 500 ft above the minimum vectoring altitude, which is determined by airport elevation, terrain clearance, and other local factors. Thus CODAS VMC corresponds to "high Visual Flight Rules (VFR)". At Boston (BOS), visual approaches are permitted when the ceiling exceeds 2500ft and the visibility exceeds 5mi. At Dallas Fort Worth Airport (DFW), visual approaches are permitted only for ceilings above 3500ft and visibility greater than 5mi. Lower ceiling and visibility conditions are considered IMC. That is, CODAS IMC corresponds to "low VFR" and below. CODAS weather data come in either 15-minute or hourly summaries. Any hour with one or more 15-minute intervals of IMC is considered to be an IMC hour. In our analysis, any day with one or more IMC hours between 6AM and midnight is considered to be an IMC day.

Using these definitions we find that at DFW in calendar year 1997, 35% of the days had one or more IMC hours between 6 am and midnight. In act the top 33 delay days were all IMC days and 38 of the top 40 delay days were IMC days. Figure 7 compares CODAS airborne delay on IMC and VMC days at DFW in 1997.

On VMC days, the mean was 1.9 minutes of delay per aircraft, the standard deviation was 0.83 minutes per aircraft, and the delay on the worst VMC day averaged 5.3 minutes of delay per aircraft. On IMC days, the means, standard deviations, and peak delays were 2 to 3 times larger than on VMC days. The CODAS arrival delay at DFW in 1997 was 3 to 4 times larger than the airborne delay by all statistical measures, and showed a similar factor of 2-3 increase in IMC.

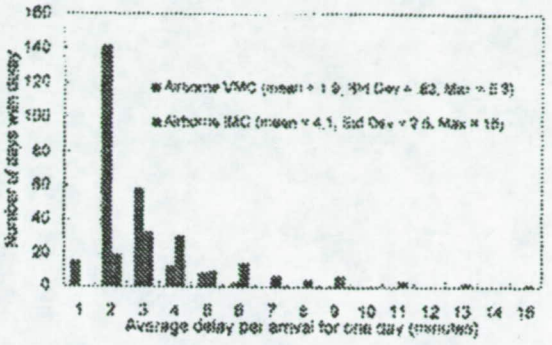

Figure 7. Distribution of CODAS airborne delay on VMC and IMC days - DFW 1997.

Figure 8 separates the CODAS airborne delay data into IMC and VMC components for 7 important airports (Atlanta (ATL), DFW, LGA, BOS, Philadelphia (PHL), Newark (EWR), and LAX) with varying operational characteristics.

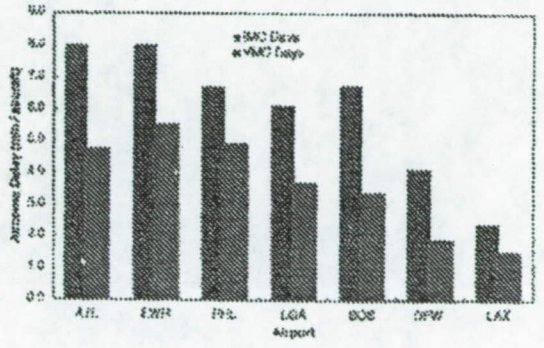

Figure 8. CODAS airborne delay for seven airports on VMC and IMC days – 1997.

Results similar to DFW were found for all of these airports: on IMC days the means, standard deviations, and peak delays were significantly larger than their values on VMC days. The observation that BOS and DFW delays were equally sensitive to IMC is somewhat unexpected. The sensitivity of arrival runway capacity to meteorological conditions differs significantly between these two airports. In some reconfiguration situations, Boston's arrival runway capacity can drop by nearly 50% in IMC, whereas the biggest IMC arrival runway capacity drop possible at DFW in 1997 was about 33%.

Although, on average, DFW has excess capacity that is not strongly influenced by reduced ceiling and visibility, DFW experiences hubbing peaks each day that temporarily approach the available VMC runway capacity. During these rushes, a small decrease in either en route or terminal capacity can cause a large increase in delay. In VMC the queues that build up in

these brief periods of excess demand are quickly cleared after the demand subsides. It takes longer to clear these queues in IMC. In 1997 before the new DFW runway became operational, arrival capacity could often be reduced by loss of a diagonal runway, resulting in even larger queues during transient arrival rushes and longer residual recovery periods after the rushes subsided.

DELAY CORRELATION

The tendency for airports to experience larger delay means and variances on IMC days than on VMC days seems to support the notion of local causality. That is, if we can increase IMC arrival capacity at an airport, we should also reduce delays at that airport. However, when we examine the correlation between delays at airport pairs, we find evidence of systematic effects correlating delays over the region for CODAS airborne delays as well as arrival delays. This occurs even on mixed days when one airport experiences some IMC and the other experiences solid VMC. We also see that correlation decreases as the geographical separation between airports increases.

Figure 9 shows the correlation between CODAS arrival delays at EWR and LGA for all days in 1997 for all four combinations of meteorological conditions. The correlation between CODAS arrival delays on the days in which IMC prevailed at both airports was 0.85. Correlation was equally strong on those few "mixed" days when it was IMC at one airport and VMC at the other. The weakest correlation was for the majority of days when the weather was clear at both airports. Even on these days the correlation was significant and positive at 0.49.

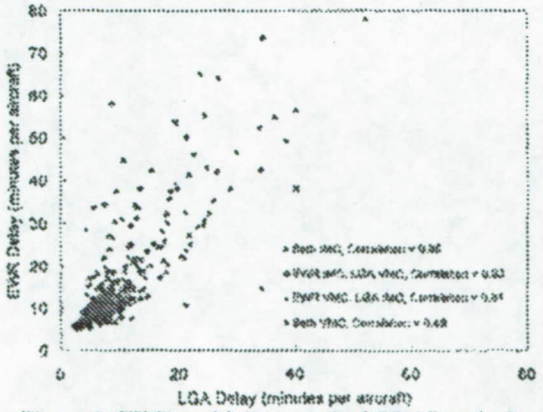

Figure 9. EWR and LGA – daily CODAS arrival delay correlation – 1997 all days.

Strong correlation between EWR and LGA delays might be expected because their traffic is managed by a common TRACON, the airports are close to each other geographically, and they share common arrival and departure fixes.[11]. Because of this physical proximity, weather conditions are also correlated between the two airports: in 1997 there were only 34 days – split 15/19 – in which one airport experienced some IMC and the other experienced solid VMC; there were 121 days when both experienced IMC; and there were 208 days when both experienced solid VMC.

Strong correlation also occurs between CODAS arrival delays at other airport pairs since delays relative to schedule are influenced by the connectivity of the air transport network. What is surprising is that strong positive correlation also occurs between airborne delays at some airport pairs.

Figure 10 shows the correlation between EWR and PHL for CODAS airborne delay. There was significant correlation even though these were not schedule-related delays, the distance between the airports is greater, and the terminal air traffic is managed by different facilities.

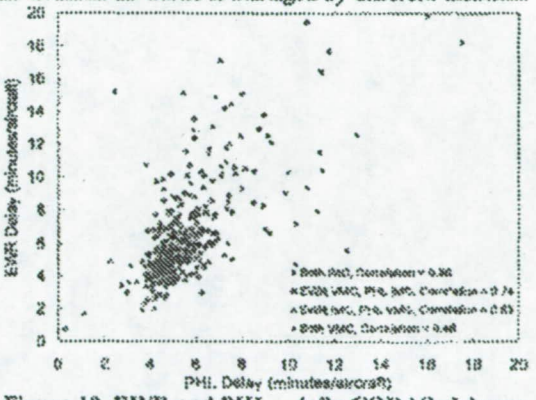

Figure 10. EWR and PHL – daily CODAS airborne delay correlation – 1997 all days.

Figure 11 summarizes the correlation coefficients between selected airport pairs for CODAS airborne and arrival delays for all days in 1997. The correlation generally decreased as the distance between the airport pairs increased.

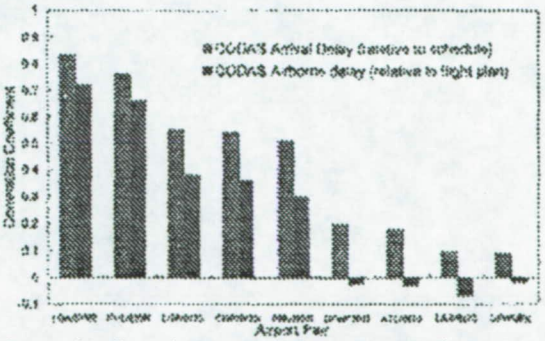

Figure 11. Correlation coefficients for CODAS delay for nine airport pairs - 1997 all days.

The annual airborne delays were not correlated for widely separated airports. However, the annual arrival delays showed small positive correlation between all of these airport pairs because of downstream schedule impacts on high delay days. For example, schedule-based arrival delays at DFW and ORD were positively correlated, probably because both airports are major hubs for American Airlines.

DOWNSTREAM DELAY

Downstream delay caused by schedule connectivity can multiply the cost of large delay events. Late arrivals propagate through airline schedules and result in additional downstream delays. This effect multiplies the dollar benefits from reductions in initial delay. We estimated the magnitude of the downstream arrival delay at DFW in 1997 based on a 1998 analysis by Beatty et al of downstream schedule delay resulting from 500 delayed departures from DFW.[12] That study showed that the number of minutes of downstream delay resulting from each initial delayed departure is related to the magnitude and the departure time of the initial delayed flight. The relationship was modeled in the form DM=1+S*DD, where DM is the delay multiplier, DD is the magnitude of the initial departure delay, and the departure-time factor, S, was obtained by a linear least-squares fit to the delay data. S is greatest when the initially delayed flight takes off early in the day. The fitted data in Beatty's final delay multiplier table shows that, to first order, S decreases linearly as the departure time of the initial delayed flight increases from 6 to 22 hours, and that S is independent of the magnitude of the initial departure delay. Figure 12 plots S versus the departure time of the initial delayed flight for three values of initial delay.

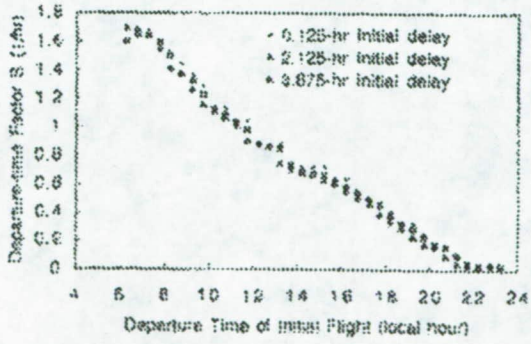

Figure 12. Departure-time factor S as a function of the time of the initial departure delay.

One can use the delay multiplier model to obtain a rough estimate of downstream delay for DFW in 1997 based on CODAS data. Because the CODAS database focuses on delay for flights *to* each airport and does not provide statistics on departure delay *from* an airport, one must use the average daily CODAS arrival delay into DFW to approximate the average daily departure delay out of DFW. This approximation ignores the first level of delay branching that occurs at DFW itself. It thus underestimates downstream delays resulting from large arrival delays and provides a lower bound on downstream delay.

We estimated the delay multiplier for each day by further assuming uniform delay between 6 and 22 hours equal to the daily average for all arrivals. The linearity of S then allows us to use the mean value of S over that period (which is 0.77/hour) to calculate downstream delay

At DFW in 1997 the daily CODAS arrival delay averages ranged from a low of about 2 minutes per flight to a high of about 58 minutes per flight. The uniform delay approximation is obviously well justified for days with very low average delay. The uniform delay approximation also appears to be a reasonable assumption for days with high average delay, although it can underestimate the downstream multiplier when a short period of very high delay occurs early in a day. Hourly arrival delay's of 3 or 4 hours occurred occasionally at DFW in 1997. However, the hourly rate only exceeded 60 minutes about one percent of the time. Most days with average daily delays exceeding 30 minutes involve episodes of relatively constant high hourly delay distributed throughout the day.

The uniform delay approximation is reasonably well justified when calculating the annual downstream delay constribution of more normal delay days. At DFW in 1997 the annual CODAS arrival delay average was about 10 minutes per flight and the annual standard deviation was about 8 minutes per flight. Although one-sigma days can experience brief periods of high delay, the occurrence times of those delays tend to be uniformly distributed when considering the large number of such days in a year.

These approximations applied to CODAS arrival delay provide a lower bound on the annual downstream delay multiplier of about 1.2 at DFW in 1997. The stacked column chart of Figure 13 shows the initial delay and the lower bound on the downstream delay estimates accumulated for all days in 1997.

The chart further breaks down delays for the 94 days in 1997 that experienced thunderstorms within 50 nautical miles of DFW. The total cumulative CODAS arrival delay for flights into DFW in 1977 was 1.43 million minutes on those thunderstorm days and 2.29 million minutes on the remaining non-thunderstorm days. Because storm days had larger initial delay, they also had larger downstream delay. Thus, the effective delay

multiplier was about 1.3 for storm days compared to 1.13 for non-storm days.

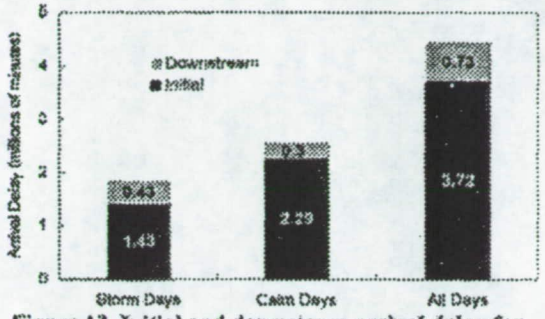

Figure 13. Initial and downstream arrival delay for days with and without storms at DFW in 1997.

HAZARDOUS WEATHER

FAST is currently unable to predict flight trajectories when storms disrupt arrival routes. Thus, thunderstorms reduce the amount of time that FAST can be used. However, such route disruptions are infrequent, and the benefit of extra runway capacity increases disproportionately when the storm has passed and controllers must clear out residual storm queues.

To determine the net effect of thunderstorms on FAST benefits it is necessary to quantify the relationship between hazardous weather and delay. We examined hazardous weather delays at DFW in 1997 and at EWR in 1999.[13] Weekly report logs from the Integrated Terminal Weather System (ITWS) at DFW indicate that there were 94 days that had thunderstorms within 50 nautical miles of DFW.[14] The DFW TRACON logs show that on about 50 of these days the storms involved enough disruption to air traffic to cause delays. At EWR there were 36 days with thunderstorms within 100 NM of the airport that caused major delays. These numbers are higher than the number of days in which thunderstorms were officially reported at DFW and EWR. Tower personnel report thunderstorms at an airport when they detect lightning or thunder. On average that occurs 45 days a year at DFW and 26 days a year at EWR.

Figure 14 is a plot of the CODAS airborne delay at DFW on the 50 worst delay days in 1997 sorted by delay magnitude. The 14 worst days all had thunderstorm activity. Thirty-four of the 40 worst airborne delay days were thunderstorm days. Large airborne delays are strongly associated with thunderstorms. Yet, because there were many more storm-free days in the year, the total annual delay on storm-free days was about 42% larger.

Figure 15 shows the cumulative 1997 CODAS airborne delay separately for thunderstorm days and all other days at DFW sorted in order of descending airborne delay. We multiplied the average delay on each day by that day's arrival count to obtain the cumulative aircraft delay minutes. The cumulative annual CODAS airborne delay on thunderstorm days was about 415,000 minutes. The cumulative delay on non-storm days was 591,000 minutes.

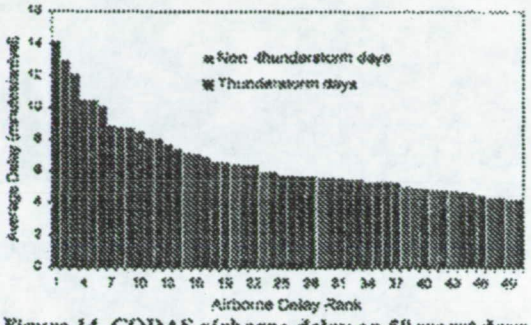

Figure 14. CODAS airborne delay on 50 worst days -DFW 1997.

The direct operating cost to airlines at DFW in 1997 can be estimated by multiplying the airborne Delays by the $19/minute estimate obtained from the Seagull and LMI benefit analyses. The results total $11.2M for non-storm days and $7.9M for thunderstorm days.

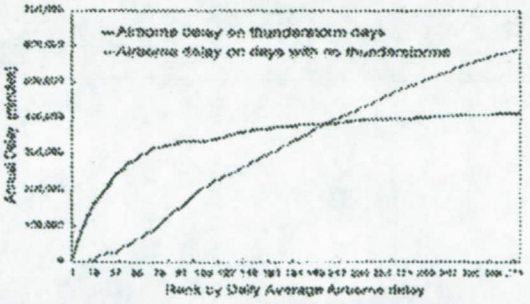

Figure 15. Cumulative CODAS airborne delay on days with and without storms– DFW 1997.

The larger cost for non-storm days occurred partly because there were more days and, to a lesser extent, because there were more arrivals on those days. The cumulative minutes and dollars for thunderstorm days would be larger if the calculation included nominal delay and dollar equivalents for each cancelled flight. As shown above, a complete cost accounting for downstream delay would also increase storm-related costs relative to costs on non-storm days because larger delays cause larger downstream ripple effects.

Thunderstorms and IMC were the main contributors to large CODAS airborne delays at DFW. But high winds alone were found to cause significant delay at EWR.[13]

Although the predominantly North-South orientation of the DFW runways makes it potentially vulnerable to crosswinds, DFW had only one day in 1997 that was free of IMC and thunderstorm activity but that had CODAS airborne delays greater than the average for an IMC day. On December 9, the delay built up during five hours of 20- to 25-kt crosswinds after 1PM, but a long period of delay in the morning when the winds were below 10 kt also contributed to the high daily delay. Unlike EWR in 1998, where winds alone caused numerous large delay events, DFW in 1997 did not experience significant delay contributions from high winds.

Delays can also be caused by inefficient handling of arrival traffic or by contention for air space and runways in peak arrival periods. But, we found that days with high average delay at DFW have statistically lower daily arrival counts. (This was seen at EWR also, where the average number of cancellations per thunderstorm or IMC event was more than 26 flights.) An airline does not cancel a flight because the demand it will generate might cause delays. Airlines cancel flights because they anticipate—or are already experiencing—costly disruptions from other causes. Although high peak demand usually increases peak delay, high daily demand is negatively correlated with high average daily delay at DFW.

Figure 16 summarizes the effect of weather on airborne delay at DFW in 1997. For this figure we initially divided the days into four categories: with and without thunderstorms and with and without periods of IMC. We found that solid VMC days with thunderstorms had mean delays almost as small as VMC days without thunderstorms, likely because the storms were far from the airport and good visibility at the airport helped clear any queues that occurred from flow disruptions. Consequently it is not necessary to distinguish between the two types of VMC days, and Figure 16 summarizes CODAS airborne delay for only three weather combinations.

As shown in part a), VMC days had the smallest average CODAS airborne delay (1.9 minutes per arrival). Days with IMC and no thunderstorms averaged 2.9 minutes of delay per arrival. Days with thunderstorms plus IMC averaged 6.1 minutes of delay by per arrival, more than double that of storm-free IMC days.

Part b) shows the number of days at DFW in 1997 that experienced each weather category. 237 days were solid VMC. 79 days had one or more hours of IMC but no thunderstorm activity within 50 NMI of the airport. And 49 days had one or more hours of IMC plus thunderstorms within 50 NMI of the airport.

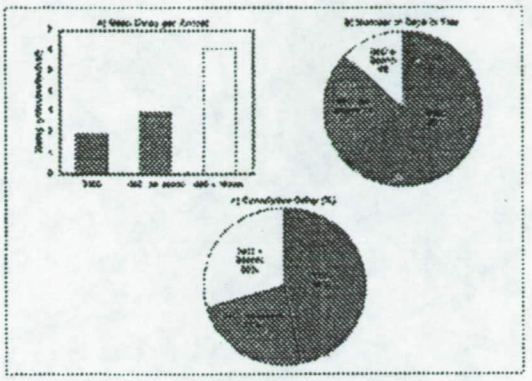

Figure 16. CODAS airborne delay statistics for three weather conditions – DFW 1997.

Part c) gives the resulting cumulative annual delay for each of the three weather conditions. Because the many small VMC delays occurred regularly during daily arrival rushes, they contributed 46% of the annual total. The 79 IMC days without thunderstorms contributed 24% of the annual total. The 49 days that had both thunderstorms and periods of IMC contributed the remaining 30%, which was the second largest cumulative annual delay. These 49 days also included 5 of the 6 ground hold days for flights into DFW in 1997.

The predominance of congestion-generated VMC delay would imply that the FAST design does indeed focus on the principal queuing problem. However, analysis of radar data at DFW has shown that a significant fraction of aircraft currently land in VMC with less than minimum radar separation.[12] FAST cannot reduce separations below radar separation standards. Thus, in VMC periods, FAST must focus principally on planning, sequencing, and runway balancing to reducing large gaps in the arrival flow.

FAST will be most effective in calm IMC periods when it can focus on achieving minimum separation standards while also reducing arrival gaps.

Storms can temporarily disrupt the operation of FAST because its algorithms currently depend on the use of standard arrival routes. FAST is thus not able to achieve any of its goals during periods when storms force aircraft to depart from standard routes. From analysis of the ITWS data logs we estimate that storms may have blocked arrival routes for as much as 50% of the day on the 49 thunderstorm days in 1997. Eliminating these disrupted periods from the modeled benefits for DFW would reduce the dollar savings from FAST by about 7 percent.

On the other hand, after a thunderstorm has passed, it often leaves a residual queue of approaching aircraft capable of landing at a rate close to the maximum IMC capacity of the airport. These are the conditions for which small incremental increases in capacity from FAST can produce very large reductions in delay. Consider for example, a situation in which the storm has left scattered aircraft that now approach from many directions at random intervals, but with an average rate that is 95% of the arrival capacity of the airport. Steady state queuing theory indicates that delay accumulates at a rate of 18 hours of delay per hour. But increasing the capacity of the airport by 5% will reduce the rate of delay accumulation by more than a factor of 2. The effectiveness of FAST is magnified significantly. As a result, the small incremental delay caused by its absence during the storm may often be offset by the delay saved with its aid after the storm.

CONCLUSIONS

Two prior FAST benefit models differ in their details and their dollar benefits for some individual airports, but predict comparable overall dollar savings. LMI modeled demand, capacity, and queuing delay more accurately, but generated FAST benefit estimates for only 10 airports. A simple steady-state model validates the LMI ranking results and makes it possible to extend the LMI dollar savings estimates to other airports. FAST benefits will accrue in all weather conditions. However, because en route and terminal airspace congestion causes queuing delay every day during arrival rushes, VMC days will contribute most of the annual delay benefit.

Benefit analyses based on the use of tighter inter-arrival spacing to reduce queuing delay tend to overestimate FAST delay benefits in VMC when FAST can only cut delay cost by reducing large gaps. Benefit analyses also tend to overestimate FAST delay benefits during those periods when normal arrival routes are totally disrupted by storms.

On the other side of the benefit ledger, the models underestimate delay benefits by ignoring reductions in downstream delay and ignoring periods following storms when large queues must be cleared. We conclude that these errors roughly balance each other so that the LMI annual savings estimates are reasonable, provided FAST can indeed reduce inter-arrival time by 15 seconds at all runways. This spacing reduction, which is equivalent to a capacity increase of 3 to 4 aircraft per hour, remains to be validated by analysis of radar data from operational tests.

REFERENCES

1. Air Transport Assocation of America, (ATA), "Industry Information - Approaching Gridlock," www.airlines.org (October 2001)
2. Davis, et al, "Operational Test Results of the Final Approach Spacing Tool," Proc. IFAC 8th Symposium on Transportation Systems '97, Chania, Greece, (June 1997)
3. Federal Aviation Administration Office of Policy and Plans (APO), "Consolidated Operations and Delay Analysis System (CODAS)," www.apo.data.faa.gov/faacodasall.HTM (October 2001)
4. Hunter, G et al, "CTAS Error Sensitivity, Fuel Efficiency, and Throughput Benefits Analysis," Seagull Technology, Inc. TR96150-02 (July 1996).
5. Couluris, G., Weidner, T., Sorensen, J., "Initial Air Traffic Management (ATM) Enhancement Potential Benefits Analysis," Seagull Technology, Inc. TR96151-01 (September 1996).
6. Weidner, T., Couluris G., Hunter, G., "CTAS Benefits Extrapolation Preliminary Analysis," Seagull Technology, Inc. TR96156-01 (February 1998).
7. Weidner, T., "Capacity Related Benefits of Proposed CNS/ATM Technologies," 2nd USA/EUROPE Air Traffic Management R&D Seminar, Orlando, FLA, (December 1998).
8. Lee et al, "Estimating the Effects of Terminal Area Productivity Program," NASA Contractor Report CR-1997-201682 (April 1997).
9. Hemm, R. et al, "Benefit Estimates of Terminal Area Productivity Program Technologies," NASA Contractor Report CR-1999-208989 (January 1999).
10. Federal Aviation Administration Air Traffic Tactical Operations Program Office (ATT-200), "Air Traffic Operations Network (OPSNET)," www.apo.data.faa.gov/op .et (October 2001).
11. DeArmon et al, "Complex, Congested Airspace and Impacts on Airline Operations," Chapter 34, p. 495, "Handbook of Airline Operations," McGraw Hill, New York, 2000.
12. Beatty, R. et al, "Preliminary Evaluation of Flight Delay Propagation Through an Airline Schedule," 2nd USA/EUROPE Air Traffic Management R&D Seminar, Orlando FLA (December 1998).
13. Allan, S., Gaddy, S., Evans, J., "Delay Causality and Reduction at Newark International Airport Using Terminal Weather Information Systems," MIT Lincoln Laboratory Project Report ATC-291, (July 2000).
14. Evans, J., Ducot, E., "The Integrated Terminal Weather System (ITWS)," Lincoln Laboratory Journal, Vol. 7, No. 2, (1994).
15. Andrews, J. W. and Robinson, J.E., "Radar-Based Analysis of the Efficiency of Runway Use," AIAA Guidance, Navigation & Control Conference, Montreal, (Aug 2001).

AIAA Guidance, Navigation & Control Conference, Montreal, Quebec, 6-9 August 2001 AIAA-2001-4361

THE DESIGN AND IMPLEMENTATION OF THE NEW CENTER/TRACON AUTOMATION SYSTEM (CTAS) WEATHER DISTRIBUTION SYSTEM*,†

Steven D. Campbell, Richard A. Hogaboom, Richard T. Lloyd, James R. Murphy‡ and Herman F. Vandevenne

M.I.T. Lincoln Laboratory
244 Wood Street
Lexington, MA 02420-9108

ABSTRACT

The National Aeronautics and Space Administration (NASA), working with the Federal Aviation Administration (FAA), is developing a suite of decision support tools, called the Center/TRACON Automation System (CTAS). CTAS tools such as the Traffic Management Advisor (TMA) and Final Approach Spacing Tool (FAST) are designed to increase the efficiency of the air traffic flow into and through Terminal airspace. A core capability of CTAS is the Trajectory Synthesis (TS) software for accurately predicting an aircraft's trajectory. In order to compute these trajectories, TS needs an efficient access mechanism for obtaining the most up-to-date and accurate winds.

The current CTAS weather access mechanism suffers from several major drawbacks.[1] First, the mechanism can only handle a winds at a single resolution (presently 40-80 km). This prevents CTAS from taking advantage of high resolution wind from sources such as the Integrated Terminal Weather System (ITWS). Second, the present weather access mechanism is memory intensive and does not extend well to higher grid resolutions. This potentially limits CTAS in taking advantage of improvements in wind resolution from sources such as the Rapid Update Cycle (RUC). Third, the present method is processing intensive and limits the ability of CTAS to handle higher traffic loads. This potentially could impact the ability of new tools such as Direct-To and Multi-Center TMA (McTMA) to deal with increased traffic loads associated with adjacent Centers.

In response to these challenges, M.I.T. Lincoln Laboratory has developed a new CTAS weather distribution (WxDist) system. There are two key elements to the new approach. First, the single wind grid is replaced with a set of nested grids for the TRACON, Center and Adjacent Center airspaces. Each and the grids are updated independently of each other. The second key element is replacement of the present interpolation scheme with a nearest-neighbor value approach. Previous studies have shown that this nearest-neighbor method does not degrade trajectory accuracy for the grid sizes under consideration.[6,8]

The new software design replaces the current implementation, known as the Weather Data Processing Daemon (WDPD), with a new approach. The Weather Server (WxServer) sends the weather grids to a Weather Client (WxClient) residing on each CTAS workstation running TS or PGUI (Planview Graphical User Interface) processes. The present point-to-point weather file distribution is replaced in the new scheme with a reliable multi-cast mechanism. This new distribution mechanism combined with data compression techniques greatly reduces network traffic compared to the present method. Other new processes combine RUC and ITWS data in a fail-soft manner to generate the multiple grids. The nearest-neighbor access method also substantially speeds up weather access. In combination with other improvements, the winds access speed is more than doubled over the original implementation

DESIGN CONSIDERATIONS

Approach

The new weather distribution design relies on replacing the current single wind grid with a set of nested wind grids and on replacing the current interpolation method with a nearest-neighbor retrieval method. These concepts will now be discussed further.

Nested Wind Grids

The nested grid approach is presented conceptually in Figure 1. The nested grids are defined for the TRACON, ARTCC and Multi-Center airspace. The

† This work was performed for the National Aeronautics and Space Administration under Air Force Contract No. F19628-00-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the author and are not necessarily endorsed by NASA.
‡ James Murphy is now with NASA Ames Research Center, MS-210-6, Moffett Field, CA.

nominal spatial resolution of the grids is 1 nm (~ 2 km) for the TRACON grid, 5 nm (~10 km) for the ARTCC grid and 20 nm (~40 km) for the Multi-Center grid. It should be noted that the grids are all aligned. That is, the grid points for all three grids can be thought of as being placed on a uniform 1 nm grid.

Each grid is rectangular in shape and sized in such a way to encompass the region of interest. That is, the TRACON grid is sized to be encompass the TRACON region plus a buffer region around it. Likewise, the ARTCC grid encompasses the ARTCC plus a buffer and the Multi-Center grid extends out into adjacent Centers a sufficient distance to allow boundary crossings to be scheduled.

These nested grids are used in the following way. For every wind retrieval, the position of the aircraft is checked to determine which of the nested grids should be used. As an example, imagine an aircraft approaching the ARTCC from an adjacent Center. The Multi-Center grid is sized to include the furthest aircraft in an adjacent Center for which a trajectory needs to be generated (e.g., to compute the Center boundary crossing time). The resolution of this grid is the same as the present single-resolution winds grid. As the aircraft comes closer to the ARTCC, the aircraft's position is checked for each retrieval to determine the appropriate nested grid (note: since the grids are rectangular in shape, this check is inexpensive computationally). When the aircraft crosses from the Multi-Center grid to the ARTCC grid, the wind retrievals are then made from the ARTCC grid. Similarly, when the aircraft enters the TRACON grid, the wind retrievals are then be made from that grid.

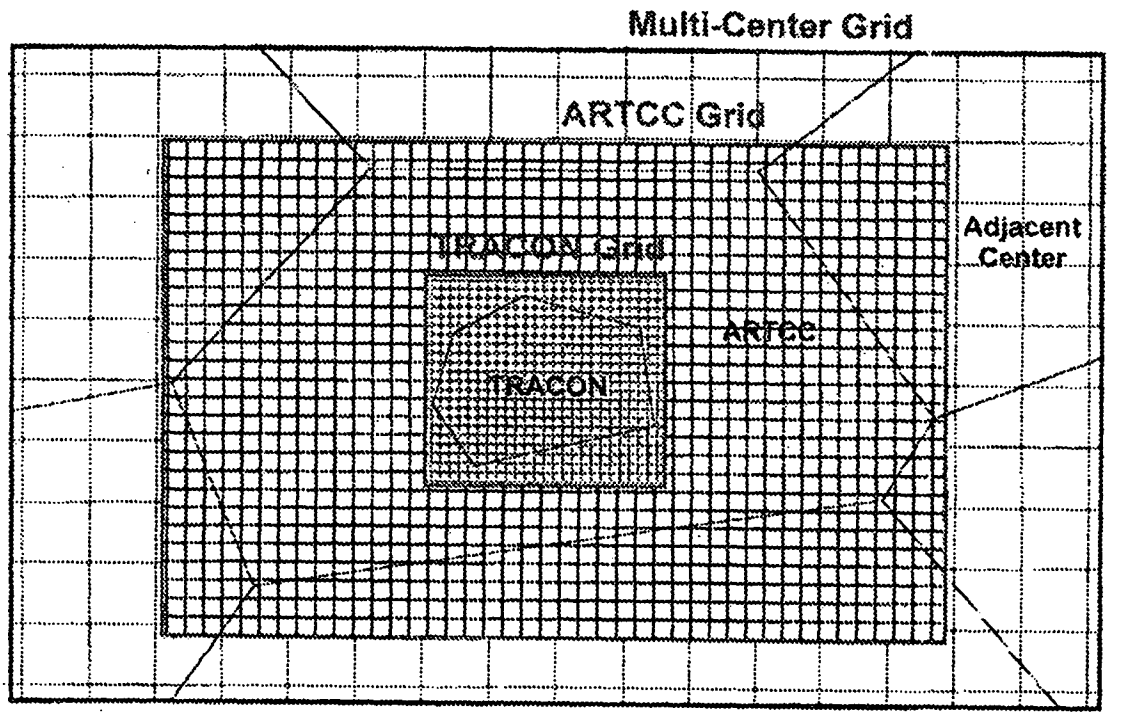

Figure 1. Nested wind grids for new weather distribution scheme.

Figure 2 shows the data sources for the three weather grids: Rapid Update Cycle (RUC) 40 km winds and Integrated Terminal Weather System (ITWS) 10 km and 2 km winds. As shown in the figure, the Multi-Center grid is generated from RUC winds alone, the ARTCC grid is generated from RUC and ITWS 10 km winds, and the TRACON grid is generated from all three sources. Moreover, all three grids can be generated from RUC winds alone if the ITWS winds are not available. Finally, there can be alternate sources for the RUC data (not shown), including multiple RUC feeds and Eta model data.

It should be noted that the domains of the weather sources and the nested grids are independent. Likewise, the update rates of the weather sources are independent of each other. When an update for a given weather source is received, the appropriate portions of the affected nested grids are updated.

Nearest-Neighbor Retrieval Method

In the previous section, it was shown how the nested grid approach allows the use of multiple wind grids updated from various weather sources with different spatial resolutions and update rates. However, a complication arises due to the greatly increased number of grid points that need to be transmitted. In the old method, an interpolation scheme was used to allow the relatively course wind grids to be accessed rapidly. However, the memory requirements for the old method were very high on a per grid point basis. Previous studies determined that the memory requirements for the interpolation method made it infeasible for extension to the nested grid approach.[4,5] Accordingly a new winds retrieval method was proposed as shown in Figure 3.

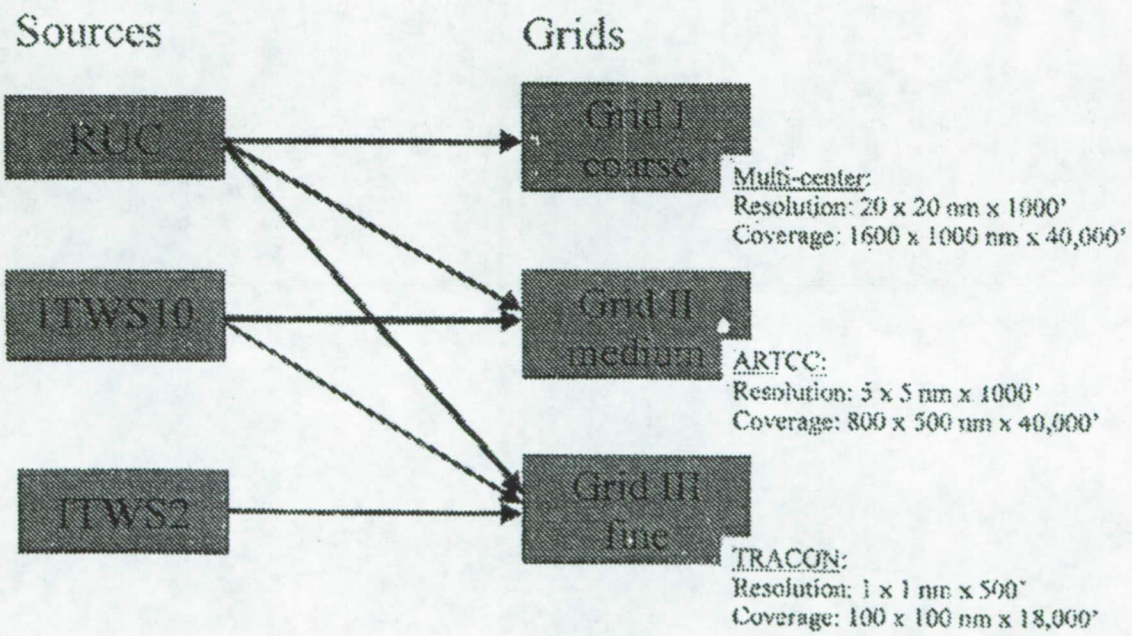

Figure 2. Data sources for weather grids.

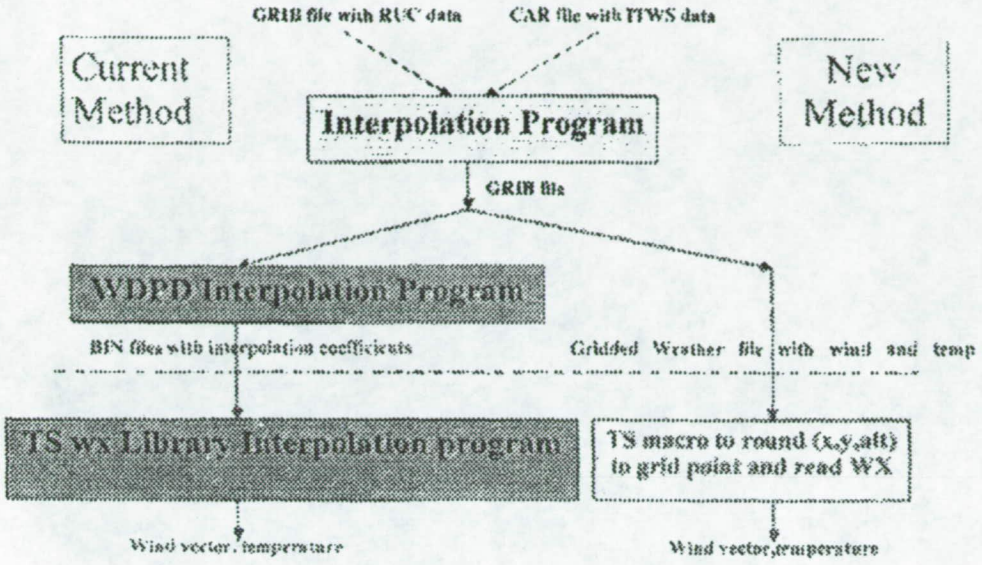

Figure 3. Illustration of replacing interpolation with nearest-neighbor retrieval method.

In the new method, the interpolation step is bypassed in favor of a nearest-neighbor retrieval scheme. The primary motivation for the nearest-neighbor approach is to make the memory requirements of the nested grids feasible, but it also has the advantage of speeding up access time and reducing processing requirements. As reported previously, the weather data access speed more than doubled with the new method (note: includes the effect of eliminating geometric altitude, which is being incorporated into the present system).[7] Figure 4 illustrates these weather data access speed improvements.

In order to employ the nearest-neighbor technique, it was necessary to verify that trajectory accuracy would not be impacted. As reported in, tests were run comparing trajectories computed using the interpolation and nearest-neighbor methods.[8] The comparison was run for two cases: a) Meter fix to threshold using ITWS 2 km winds and b) Coordination fix to meter fix using RUC interpolated to 10 km & ITWS 10 km winds.

It was found that the use of the nearest-neighbor method produced a one second RMS difference in trajectories for the first case and a four second RMS difference for the second case. These differences are negligible for the trajectories examined. An example of these results for ITWS 2 km data is shown in Figure 5. Note: the effect of using the nearest-neighbor method for the Multi-Center grid was not examined, however, the effect is assumed to be insignificant given the large distances involved and low update rate of the adjacent Center traffic data (e.g., 3 minutes via ETMS). A recent study showed the effect of using nearest-neighbor access for 40 nm winds for trajectories in the Center was a ten second RMS difference vs. interpolation.[9]

SOFTWARE DESIGN

Overview

A block diagram of the Weather Distribution (WxDist) software is shown in Figure 6. The key modules are the Weather Server (WxServer) and Weather Client (WxClient) modules. The WxServer module provides the weather data to multiple WxClient modules. There is on WxServer module for a given CTAS installation, and there is one WxClient module for each workstation employing one or more TS or PGUI processes. The external weather sources are converted into GRIB (Gridded Binary) files and divided up into minor grids as described in the previous section. The minor grids are transmitted to the WxClient processes via the reliable multi-cast protocol over the local area network (LAN). The WxClient processes provide the weather data to the application processes via a shared memory interface. The Weather Library (Wx Library) accesses the shared memory and provides the interface to the weather users.

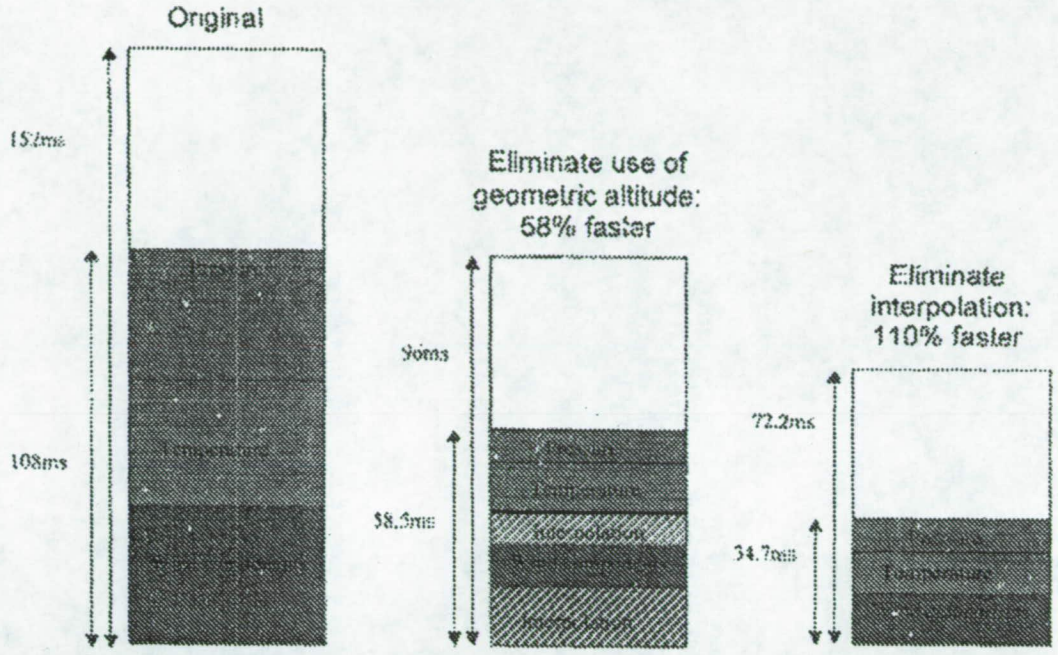

Figure 4. Weather data access speed improvement.

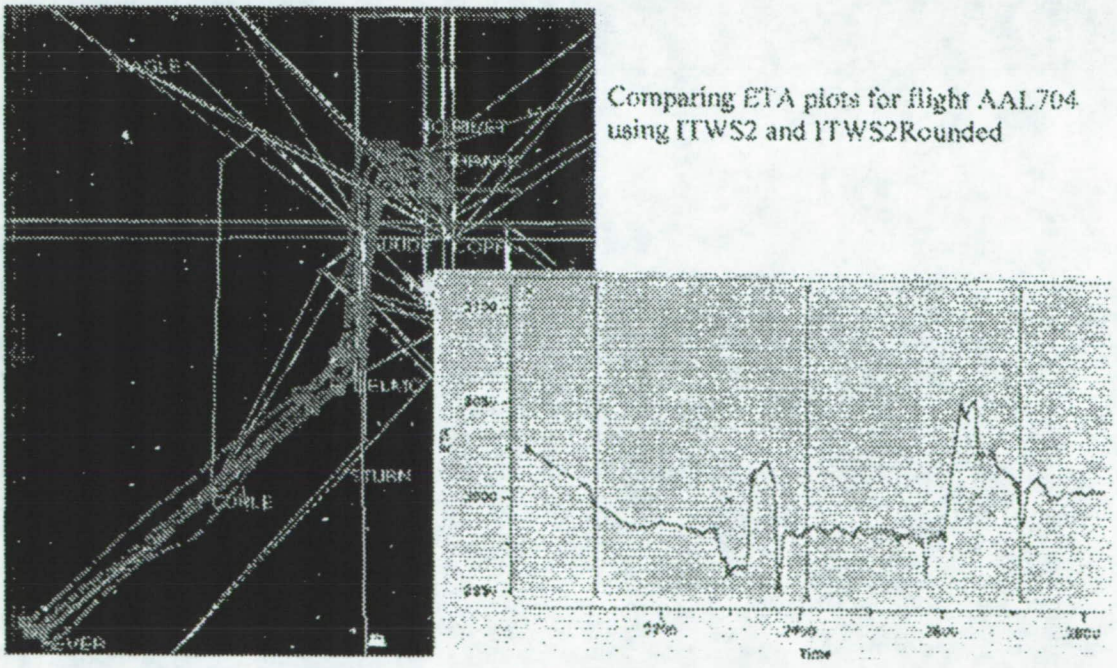

Figure 5. Comparison of trajectories for interpolated vs. nearest-neighbor methods.

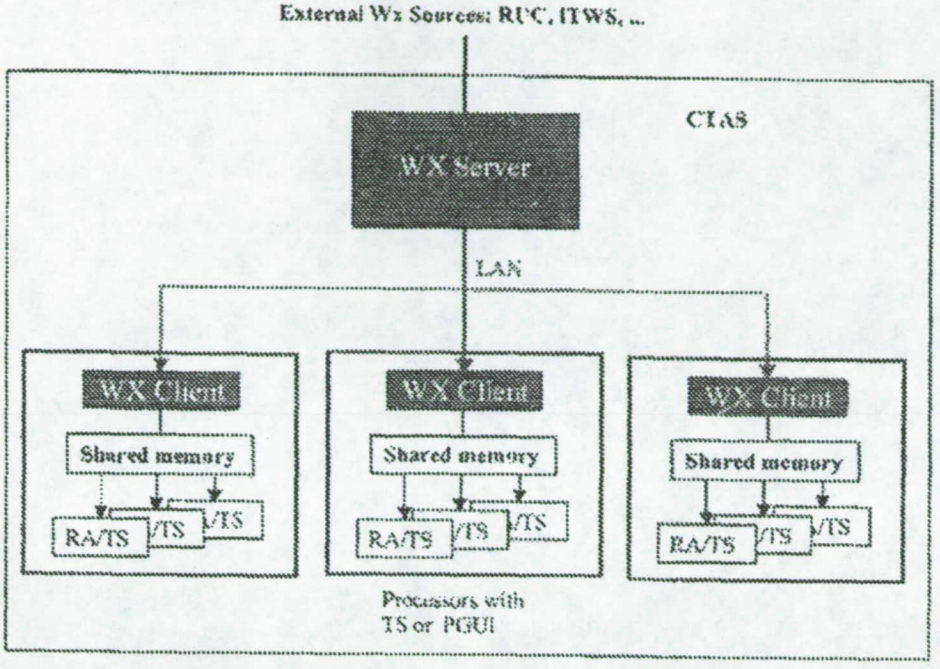

Figure 6. Weather distribution system block diagram.

Figure 7 shows a more detailed view of the CTAS weather distribution modules. As shown in the figure, the ITWS data is processed by the ITWS Connection Module into GRIB files and passed to the WxServer. The WxServer module also receives RUC files in GRIB format from the existing WDAD (Weather Data Acquisition Daemon) process. Note that the processes for acquiring the ITWS and RUC data reside outside the firewall to isolate the weather sources from the CTAS system. The CTAS Weather Communication module implements the reliable multi-cast protocol for transferring the minor grids over the LAN. Finally, the WxLibrary module provides the interface between the weather distribution system and the applications.

Table 1 lists the major modules in the new weather distribution system. The functionality of each module is summarized in the table. These modules will now be briefly described.

ITWS Communication Module

The ITWS Communication Module (itws_rtdc) obtains the Terminal Winds data feed from the ITWS testbed via stream connections and converts the data from Cartesian (CAR) format to GRIB (Gridded Binary) format. It establishes socket connections to transmit the converted ITWS Terminal Winds data to multiple local and/or remote Weather Servers.

Weather Server

The Weather Server (WxServer) module sets up the site-adaptable wind grids from a WxServer configuration file. The file defines the resolution, spatial extent and data sources for each grid. It also defines the backup strategy for generating these grids from alternate sources.

The WxServer accepts Terminal Winds data from the ITWS Communications module via socket connections and RUC data from the existing Weather Data Acquisition Daemon (WDAD) process via file transfer. The WxServer automatically selects between the alternate weather sources to determine the best available data. It then translates the input data to the NAS coordinate system and compresses the data for transmission using the GRIB compression algorithm. It then transmits the data to each Weather Client using the reliable multicast protocol.

The details of the automatic source selection logic is shown in Figure 8. In this example, the ITWS 2 km winds data feed is interrupted at time T_1. After a timeout period (nominally 6 minutes), this weather source is declared unavailable. However, the last data received continues to be used until the nominal ITWS 10 km update time.

If the ITWS 10 km update is received, the ITWS 2 km grids are now generated from the 10 km winds data. If the 10 km data is not available, then this source is declared unavailable and the most recent data continues to be used until the nominal RUC update time. When the RUC update is received, then the ITWS 2 km and 10 km grids are now generated from the RUC data. If the ITWS data later becomes available, then the system returns to using that data to generate the 2 km and 10 km wind grids.

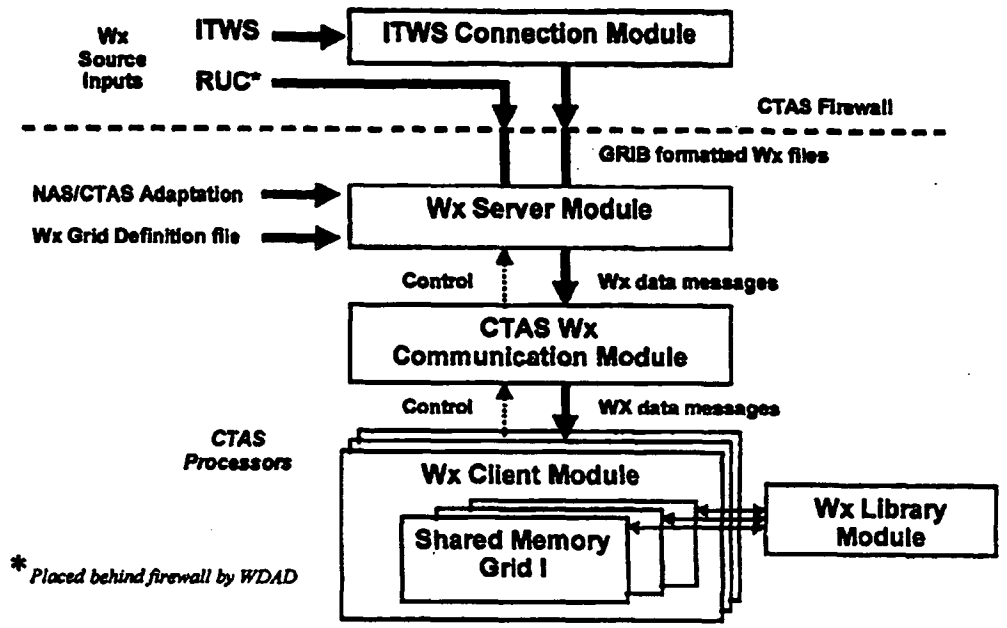

Figure 7. CTAS weather distribution modules.

Table 1. Major Software Modules

NAME	FUNCTIONALITY
ITWS Connection Module	• Connection to ITWS for winds data • Data conversion to GRIB format • Connection to multiple WxServers
WxServer Module	• Site adaptable weather grids • File connection for ITWS data • Conversion to NAS coordinates • Automatic weather source selection • Weather data compression
Communications Protocol	• Automatic connection/reconnection • Reliable multicast to transmit compressed weather data
WxClient Module	• Read multicast weather data • Update shared memory buffers • Switch buffers on command
WxLibrary Module	• When user asks for weather products: 1. Determine appropriate grid 2. Select nearest-neighbor value 3. Read weather product • Switch memory page on command • Identify current weather sources Note: grid structures transparent to users.

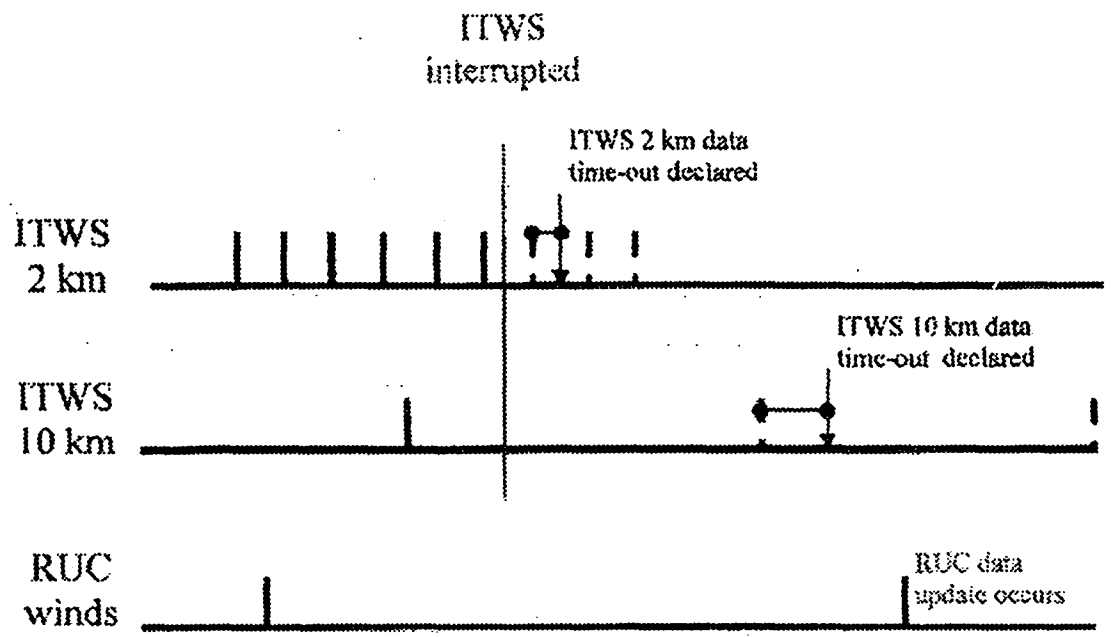

Figure 8. Timeliness check logic illustration.

WxServer/WxClient Communications Protocol

Figure 9 summarizes the WxServer/WxClient communications protocol. The WxServer process delivers the nested grid data to one or more WxClient processes using a reliable multicast protocol. This protocol makes use of a TCP/IP based socket connection between the WxServer and each WxClient as a control channel. It uses IP Multicast as the data channel from the WxServer to all WxClients.

The use of TCP/IP for the control channel provides reliability and ensures that no control message will be lost or delivered out of sequence. This guarantees the stability of the shared data view within the WxServer/WxClients group. The use of IP Multicast for the data channel allows us to transmit the data once for all WxClients regardless of how many are in use on the CTAS system. This minimizes the network load associated with weather data and allows the transmission of higher spatial and temporal resolution data sets.

WxClient

The Weather Client (WxClient) reads the weather grids via the reliable multicast protocol and updates the appropriate area of shared memory. The details of the shared memory interface with the weather users is discussed below.

Shared Memory Interface

Shared memory is used by the WxClient to make the weather data available for the Wx User. Two buffers are used for each data grid, one contains the current data and is available for reading, the other is being written to with new data. The form of each buffer is fixed with a standard header followed by a three dimensional array of product structures.

The double buffered approach used for shared memory is illustrated in Figure 10. As shown in the figure, new data is written into the write page while the users access the read page. When the update is complete, the pages are swapped. There is a signal handling scheme employed to ensure that all the users have the current information before the swap is carried out. Another implementation detail is that the contents of the new read page need to be copied back to the old read page prior to allowing further updates to occur to the new write page.

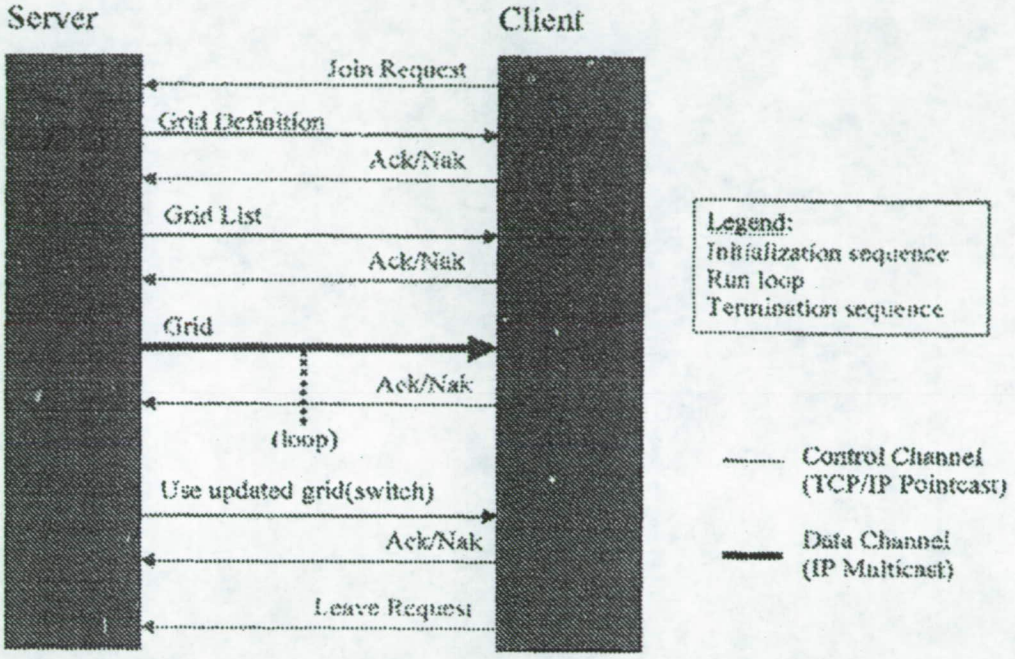

Figure 9. WxServer/WxClient Communications Protocol.

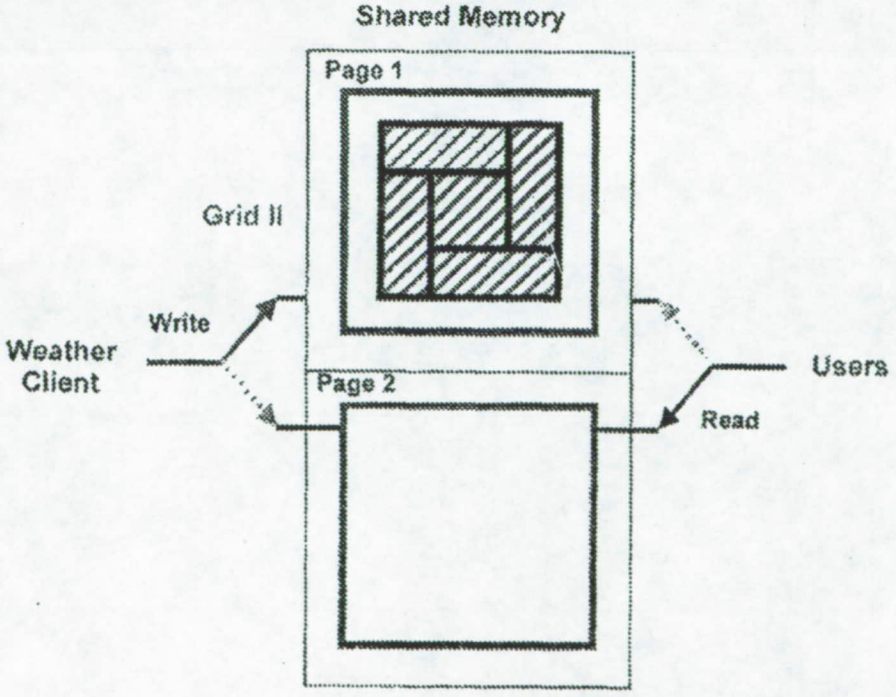

Figure 10. Shared memory double buffered scheme.

Support for Multiple Weather Users

Each WxClient is capable of supporting multiple weather users. Recall that there is one WxClient for each workstation that supports all the weather users residing on that physical piece of hardware. As described in the previous section, a double buffer scheme is used to allow the weather users to read the current weather data from one buffer while the other buffer is being filled with the next weather data to be used. An important consideration is to ensure that all weather users switch from one buffer to the other buffer in concert. This synchronization is carried out using a semaphore mechanism.

WxLibrary

The Weather Library API provides the necessary interface to the weather server for all weather using applications. The intent is to allow the weather using application to request one or more weather products at one or more locations in a single request. Units used for both the locations and the weather products are felt to be those most suitable to the using application.

ALGORITHMIC ISSUES

There are three issues which have been identified as potentially requiring changes to TS in order to accommodate the new weather distribution system. These issues include:

- gradient computation,
- temperature interpolation and
- capture condition completion.

These issues will now be discussed.

Gradient Computation

The first TS weather use issue is the calculation of the vertical wind gradient. This gradient was originally calculated by accessing wind data at two slightly different altitudes h and h+Δh, taking the difference and dividing by Δh to obtain a discrete approximation of the vertical wind gradient.

In the original system, a value of 50' was used for Δh, which works properly when interpolation is used. However, when nearest-neighbor retrieval is used with this small value, the coordinates h and h+Δh usually round off to the same grid point and produce the same wind values. The resulting gradient is then zero.

In an exceptional case, the second value would round to the next altitude level and the resulting gradient would be very large. The TS software copes with large gradients by limiting the vertical wind gradient to 10 knots. The zero gradient simply means that the effect of wind is ignored. In neither case does TS fail in its trajectory calculations.

One specific instance where the vertical gradient is used is in the en route portion of the trajectory: specifically the Constant CAS and Constant Mach segments. As part of the TS algorithm, a system of differential equations is solved using a discrete step method called Runge-Kutta. This applies in particular calculations of True Airspeed, vertical speed and ground speed.

The principal equation in TS in solving for V_t is (Eq 1):

$$\frac{dV_t}{dt} = \frac{T - D}{m} - \frac{w}{m}\gamma_a + \frac{d}{dt}(V_w \cos\theta_{rw}) \quad (1)$$

where γ_a is the aerodynamic flight path angle, T is thrust, D is drag force, m is aircraft mass and w its weight, V_w is the wind speed and θ_{rw} is the relative wind angle.

For constant Mach or constant CAS, the above differential equation reduces to an algebraic expression. This is apparent if one considers the fundamental relationship (Eq. 2):

$$V_t = a(h)M(V_{CAS}, h) \quad (2)$$

where $a(h)$ is the speed of sound (as a function of altitude) and $M(.,.)$ is the Mach number (as a function of V_{cas} and altitude). It is clear that when Mach or CAS is constant, that the true airspeed is a simple function of altitude.

One can then write (Eq. 3):

$$\frac{dV_t}{dt} = \frac{dh}{dt}\cdot\frac{dV_t}{dh} \quad (3)$$

where (Eq. 4)

$$\frac{dh}{dt} = V_t \cdot \gamma_a \quad (4)$$

Making the approximation that the wind velocity over short x or y distances is constant (i.e. only dependent on h) then (Eq. 5):

$$\frac{d}{dt}(V_w \cos\theta_{rw}) = V_t.\gamma_a.\frac{d}{dh}(V_w.\cos\theta_{rw}) \qquad (5)$$

This then leads to the expression for the aerodynamic glide slope (Eq. 6):

$$\gamma_a = \frac{(T - D).g}{w[V_t(\frac{dV_t}{dh} - \frac{d}{dh}\{V_w.\cos\theta_{rw}\})]} \qquad (6)$$

Note that the wind gradient shows up in the denominator. Thus, the reason for limiting the gradient value to a maximum value is to avoid a singularity condition for the aerodynamic glide slope and eventually a sign reversal. This would certainly make TS fail, since in a descent phase it then could not satisfy the required boundary conditions. It can also be seen that if the wind gradient comes out to be zero, the only effect is that the glide slope used in the trajectory prediction is a slightly small.

But the same formula uses weather in other ways: the altitude derivative of the True Airspeed uses temperature readings at two different altitudes. This would lead to problems also if the chosen altitude difference is a small value.

The solution for both the wind gradient and the True Airspeed gradient is to force the altitude increment to be equal to the weather grid vertical spacing. In the case of nested weather grids that increment could be different for the different nested grids. This introduces some complications in the event that the gradient computation crosses the nested grid boundaries. For this reason, it would be preferable to introduce an explicit call for the gradient at a particular point. This would allow these complications to be isolated from the weather user.

Temperature Interpolation

The second TS weather use issue involves temperature interpolation. One important TS function is to meet capture conditions, such as matching cruise and descent segments to identify the top of descent (TOD). In certain cases involving iteration to meet capture conditions, the TS was found to fail when nearest-neighbor retrieval was used. This is because TS uses small changes in temperature with altitude to drive the solution in the correct direction. When the nearest-neighbor value is used, the TS could possibly to converge because the same temperature value is always retrieved.

This problem was observed for the ARTCC grid when a grid spacing in altitude of 2000' was used. However, it was found when the altitude grid spacing was reduced to 1000', no TS failures were observed. As previously noted, the grid spacing in altitude can be traded against the horizontal grid spacing without increasing the memory required if even greater vertical resolution is needed (e.g., the vertical grid spacing could be decreased from 1000' to 250' if the horizontal spacing was increased from 10 nm to 20 nm).

Another considered was a fix to always linearly interpolate the temperature values between the nearest-neighbor values above and below the current altitude. This approach incurs a minor performance penalty due to the need to retrieve two temperature values (instead of one) and to perform a simple interpolation. However, this approach has the virtue of keeping the API unchanged. It also improves the accuracy of the temperature values, which feature a strong dependence on altitude. However, in practice it has been found that the decrease in vertical grid spacing proved sufficient and this fix was not implemented.

Capture Condition Completion

A problem was found in the way that TS completes the capture condition iteration. When TS iterates to point where the altitude is within the capture limits, it stops and returns the capture altitude. For example, the desired capture altitude might be 25,000' and the actual capture altitude might be 25,010'. Originally, TS did not recompute the derived variables (CAS, ground speed, etc) for the desired capture altitude but returned the derived variables for the actual capture altitude instead.

This behavior caused problems in the TS computations due to the nearest-neighbor retrieval for temperature. Fore example, if the gridded winds layers had 2000' vertical spacing, then the temperature might be 433 °R at 24,000' and 426 °R at 26,000'. The nearest-neighbor temperature value would therefore be 433 °R at 25,000' and 426 °R at 25,010'. This would create a discontinuity in the temperature values between flight segments that might cause TS to fail.

A fix was implemented is to force TS to recompute the derived values at the end of the iteration for the desired (instead of actual) capture altitude. This fix has been accepted for incorporation into the CTAS baseline software. It should be noted that the proposed temperature interpolation fix would also address this particular problem.

TESTING PROCEDURES

Tests were carried out on the prototype implementation to validated functionality and measure performance. Note: the results presented here should be considered preliminary in nature and subject to further refinement.

Functionality Testing

The functionality tests were carried out in several steps. The first step was to perform a regression test using RUC data only. The second step was to add ITWS winds and validate correct insertion into the wind grids. The third step was to validate that the wind grids continue to be properly generated when the ITWS winds are transiently added and removed. The fourth step was to verify proper TS operation in the presence of wind field discontinuities. The fifth step was to quantify the difference in ETA (Estimated Time of Arrival) values with the addition of ITWS winds.

For testing purposes, a version of the new system was created which returns the linearly interpolated weather value instead of the nearest-neighbor value. This version (WxDist Interpolated) is not intended for operational use (since it runs more slowly than the nearest-neighbor version) but allows direct comparison between the weather values returned by the new vs. old systems from the getWeatherValue function.

The CTAS software was also instrumented to generate ETA logs, ETA log summaries, track logs and getWeatherValue logs. The standard output is to produce the ETA log, ETA log summary and track log (if radar track data is used). If verbose output is selected, then the getWeatherValue log is also produced (generally limited to short runs due to the large volume of output generated).

A capability to generate synthetic RUC data sets was also implemented to assist in regression testing. A utility program was written which allows synthetic RUC data to be generated for various test conditions. For example, one file was created with uniform wind values and temperature values that increased linearly with pressure level, and a second file was created with uniform temperature at all levels and U &V values that increased linearly with RUC X & Y coordinates, respectively. These files were used to validate the RUC-to-CTAS coordinate transformation in the vertical and horizontal dimensions, respectively.

Additional utility programs & Unix scripts were written for examining the input RUC data and processing the output test data. Unix scripts were also written to simplify making test runs and to automatically save the output logs. Processing and examination of the test data was primarily done using IDL and Excel.

Regression Testing

Regression testing was carried out to ensure that the new weather distribution system preserves the CTAS functionality. In particular, it is necessary to ensure that trajectories are correctly generated with the new vs. old systems. In order to do this, it is necessary to use RUC data only, since the old system cannot ingest ITWS winds data.

RUC winds ingestion

The first regression test was to verify that the RUC data is properly ingested into the new system. In order to do this, actual and synthetic RUC files were input to the WDPD, WxDist Gridded (nearest-neighbor) and WxDist Interpolated versions of CTAS. The outputs of the three versions were then compared and any inconsistencies diagnosed. In the process of carrying out the regression testing, several problems were found in the WDPD processing which were diagnosed and fixed. These changes are being evaluated for incorporation into the current CTAS baseline software but will not be further discussed here.

Examples of this comparison are shown in Figures 11 and 12. For these examples, the getWeatherValue calls were logged using the verbose option for a single aircraft trajectory. Figure 11 shows the retrieved temperature vs. altitude and Figure 12 shows the retrieved U wind vs. altitude. As seen in the figures, the WxDist gridded (nearest-neighbor) retrievals exhibit the expected staircase behavior whereas the WDPD and WxDist interpolated values vary smoothly. (Note: there is a small anomaly in the WDPD results below 5000' which is currently being investigated.)

ETA comparison

The second regression test was to compare trajectories generated using the WDPD and WxDist Gridded weather distribution schemes. For this test, 489 flight plan trajectories were compared for 158 aircraft using DFWF RUC weather data from November 22, 2000 at 1400Z. Figure 13 plots the ETA difference between the new and old systems for meter fix to threshold trajectories. The worst case differences range from -7 s to + 15 s. The mean ETA difference was 0.9 s (0.12%) and the RMS ETA difference was 3.3 s (0.44%). Also shown is the comparison for the WxDist Interpolated version which produced identical results to the WxDist Gridded version.

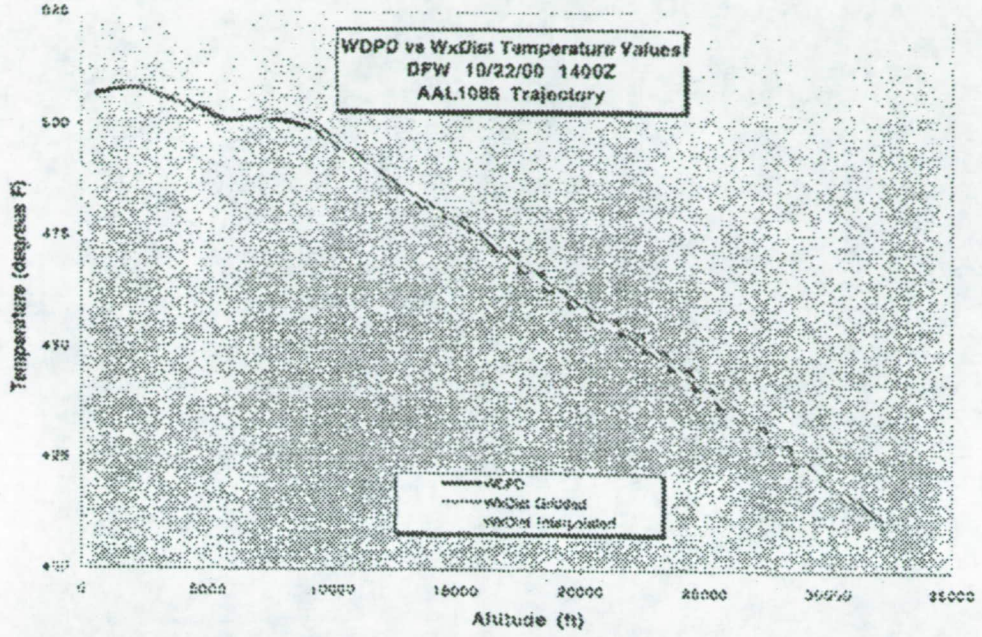

Figure 11. Comparison of retrieved temperature values for new (WxDist) vs. old (WDPD) systems

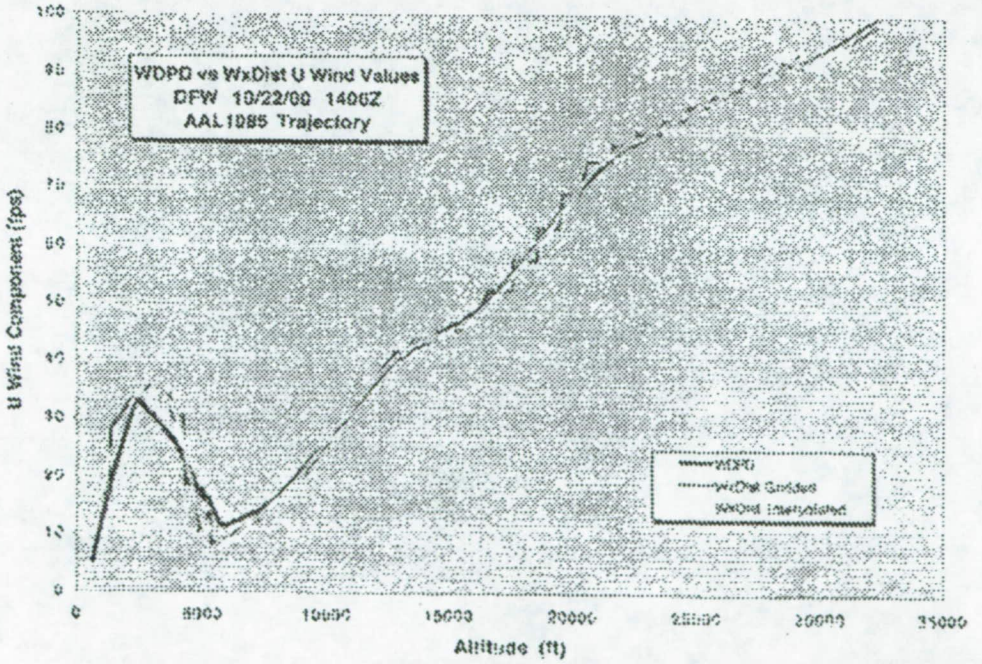

Figure 12. Comparison of retrieved U wind values for new (WxDist) vs. old (WDPD) systems.

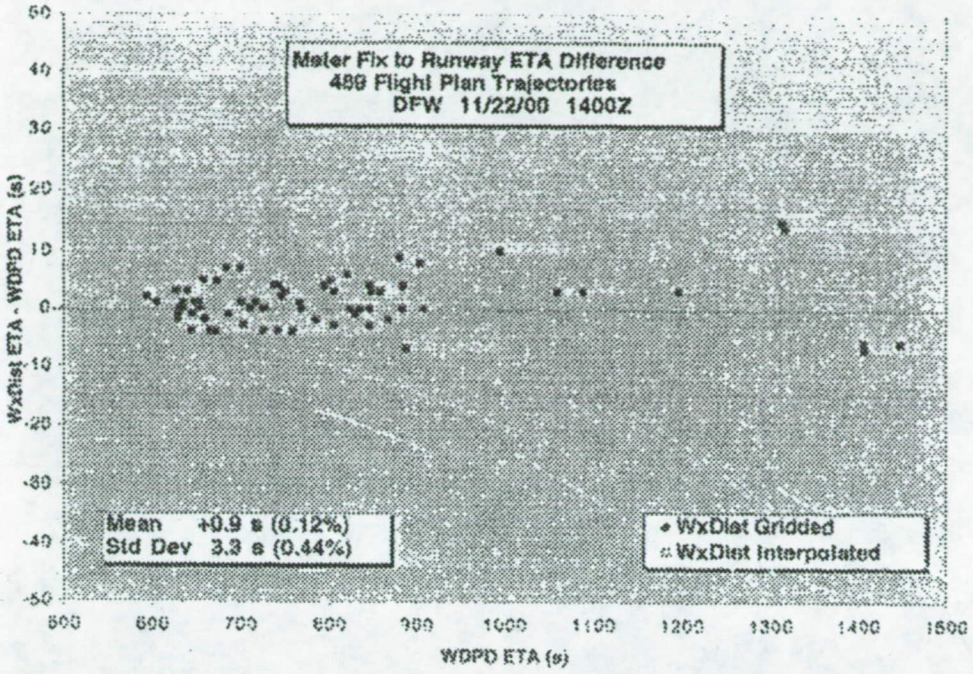

Figure 13. Flight Plan ETA difference between WDPD and WxDist for meter fix to threshold trajectories.

Figure 14 shows the ETA difference for the same set of aircraft for trajectories from the coordination fix to the meter fix. The worst-case ETA differences range from -45 s to +31 s. The mean ETA difference was -1.0 s (0.06%) and the RMS difference was 5.6 s (0.33%). Generally, the values matched very closely with the exception of a few large differences which are currently under investigation. Also shown is the comparison for the WxDist Interpolated version which were essentially the same as the WxDist Gridded results except that there was one fewer large excursion.

These error differences are slightly larger than predicted by reference (6). That study predicted ETA differences of 1 s RMS for the meter fix to threshold case and 4 s RMS for the coordination fix to meter fix case based on the use of nearest-neighbor vs. interpolated weather value retrieval. However, the ETA differences between the WxDist Gridded and WxDist Interpolated versions are in fact much smaller that predicted by the study. The observed variation is therefore likely to be due to as yet undiagnosed differences between WDPD and WxDist. However, the test results show that these differences have been reduced to very small values.

A comparison between the new and old systems was also made for aircraft trajectories from radar tracks. Figure 15 shows the Time To Fly (TTF) from 426 radar track trajectories for an aircraft landing on runway 17L. For this case, it should be noted that the WxDist ETA values were identical for the gridded and interpolated versions. The worst-case differences in TTF (ETA-current time) were -7 s to +5 seconds. The mean difference was 1.1 s (0.23%) and the RMS difference was 1.5 s (0.35%). Additional regression testing is in progress, but these results suggest that the new and old systems are working nearly identically for radar tracks.

ITWS Winds Validation

The next testing step was to validate that ITWS 10 km and 2 km winds are inserted correctly into the wind grids. For this test, the wind grids were defined as shown in Figure 16. For these tests, separate major grids were used for the ITWS 10 km and 2 km data. These major grids were set to correspond to the maximum extent of the ITWS wind fields.

In order to readily verify proper insertion of the ITWS winds, a synthetic GRIB data tool was used to generate dummy ITWS wind files. Two files were created, one for ITWS 10 km with uniform winds due East and the other for ITWS 2 km with uniform winds due North. These ITWS wind files were then processed by the Weather Server with the RUC wind file for 23 August 2000 at 1800Z for 5200' (850 mb) to produce the multiple wind grids.

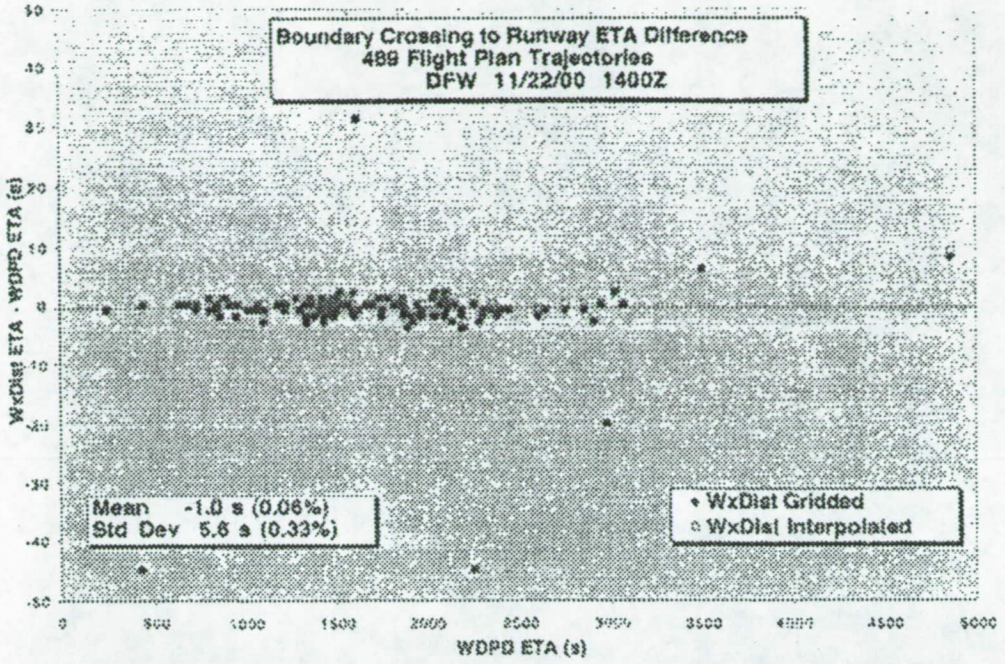

Figure 14. Flight Plan ETA difference between WDPD and WxDist for coordination fix to meter fix trajectories.

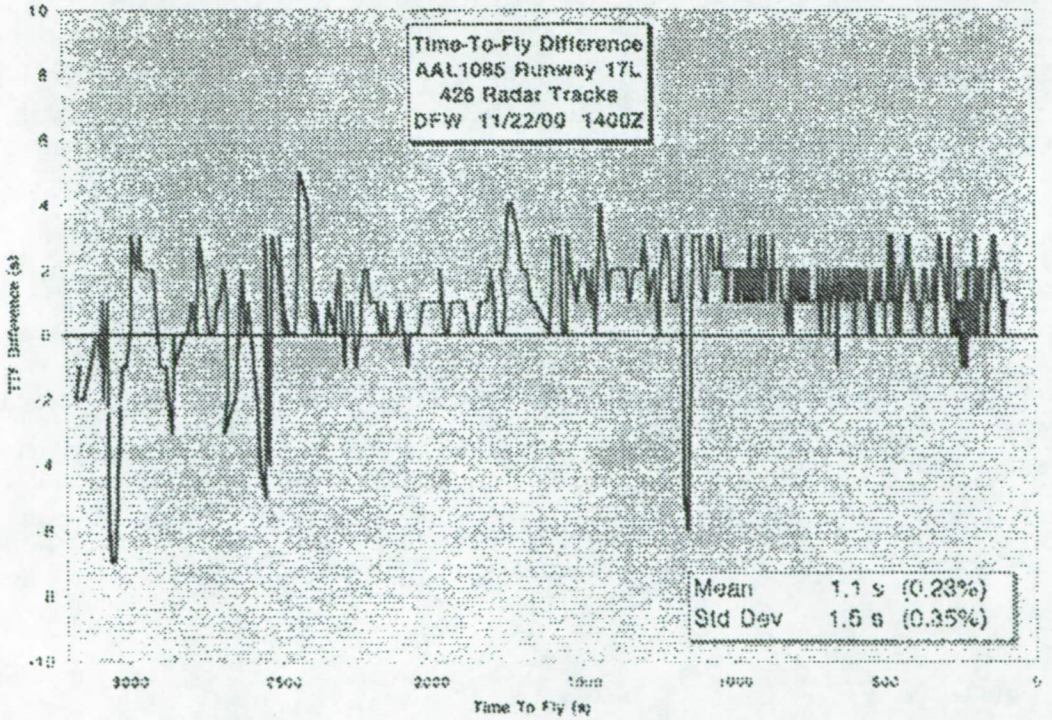

Figure 15. Time To Fly (TTF) difference for aircraft trajectories from radar tracks.

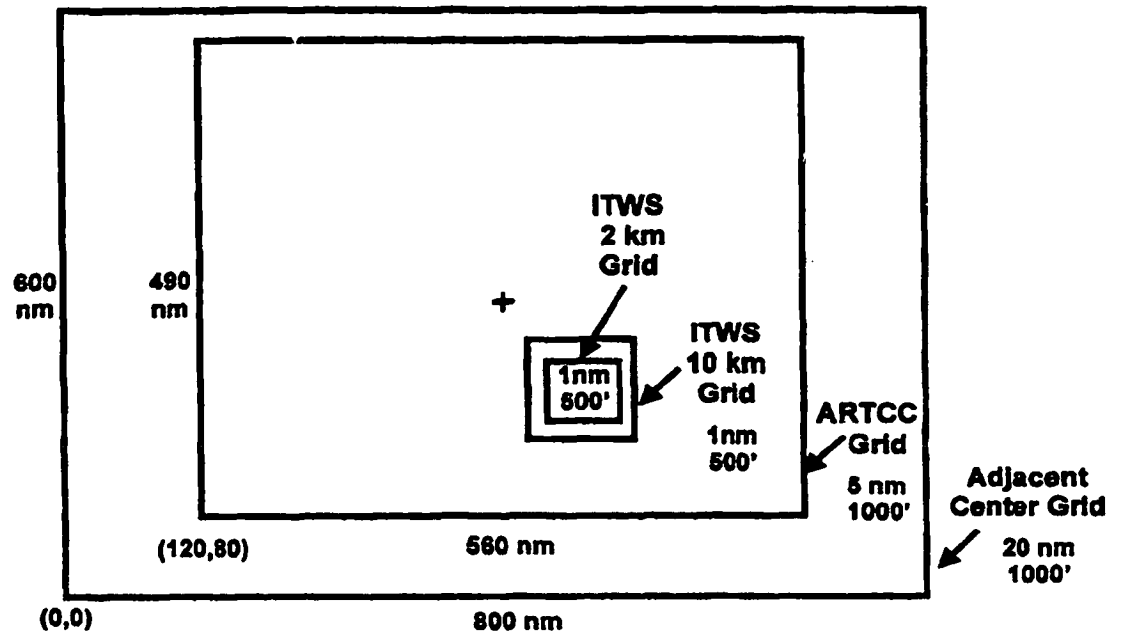

Figure 16. Major grid definitions for ITWS winds integration test.

The wind field values were retrieved on a grid approximating the 40 km RUC grid. These results are as shown in Figure 17. As seen in the figure, the insertion of the dummy ITWS wind data is clearly demonstrated.

Transient ITWS Winds Availability Testing

The third functionality test was to verify proper operation with transient ITWS winds availability. This test was carried out by interrupting the ITWS winds availability during a normal run and verifying that all the major grids continued to be generated from the RUC winds only.

Wind Field Discontinuity Testing

The fourth functionality test was to verify proper TS operation in the face of wind field discontinuities between the RUC and ITWS data. In order to rule out this possibility, CTAS was run with the dummy ITWS files shown in Figure 17. The results were examined and showed no evidence of TS failures in the face of worst-case discontinuities.

Effect on ETAs of Including ITWS Winds

The fifth functionality test was to quantify the effect on ETAs of including ITWS winds, as shown in Figure 18. To carry out this test, CTAS was run with RUC-only winds vs. RUC plus ITWS winds. Earlier studies showed considerable variation in meter fix to threshold ETAs as a function of time should be observed with the inclusion of ITWS 2 km winds updated every five minutes.[5,7]

Figure 18 summarizes the result of computing 489 pFAST (Passive Final Approach Spacing Tool) trajectories from 158 arrival flight plans under two conditions: 1) RUC winds only with interpolated winds (old system) and 2) RUC + ITWS winds with gridded winds (new system). These results were computed for DFW on 11/22/00 from 1500Z to 1700Z with RUC winds updated hourly, ITWS 10 km winds updated every 30 minutes and ITWS 2 km updated every 5 minutes.

As seen in the figure, the ETA values for the RUC only winds increase by 3 seconds over the two hour period. By contrast, the ETAs for the RUC + ITWS winds change by 16 seconds over this time period. These results are consistent with the earlier studies.

Performance Testing

Performance testing was carried out to assess the processing speed and memory requirements for the new vs. old systems. Table 2 shows examples of preliminary results for various grid sizes.

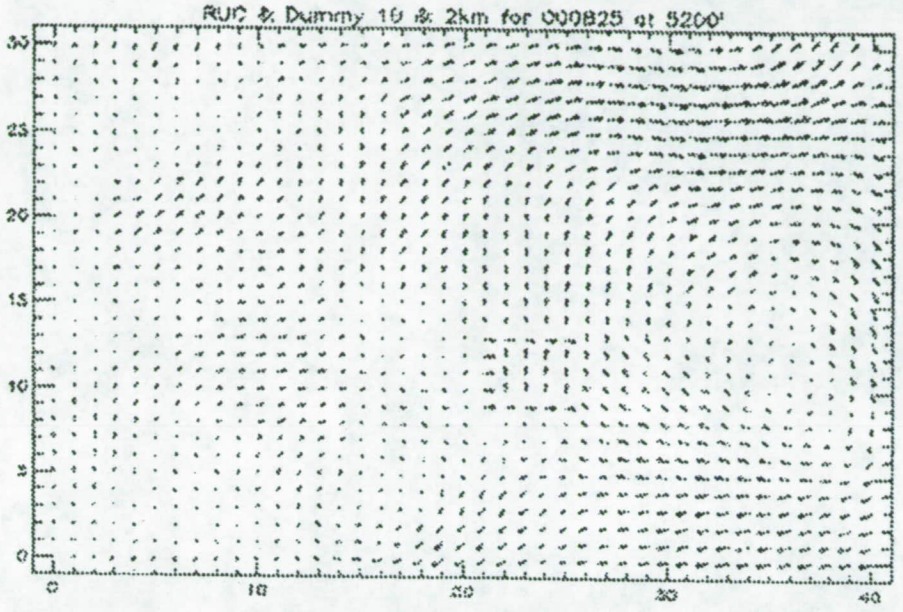

Figure 17. Example of ingesting dummy ITWS 10 km and 2 km wind field.

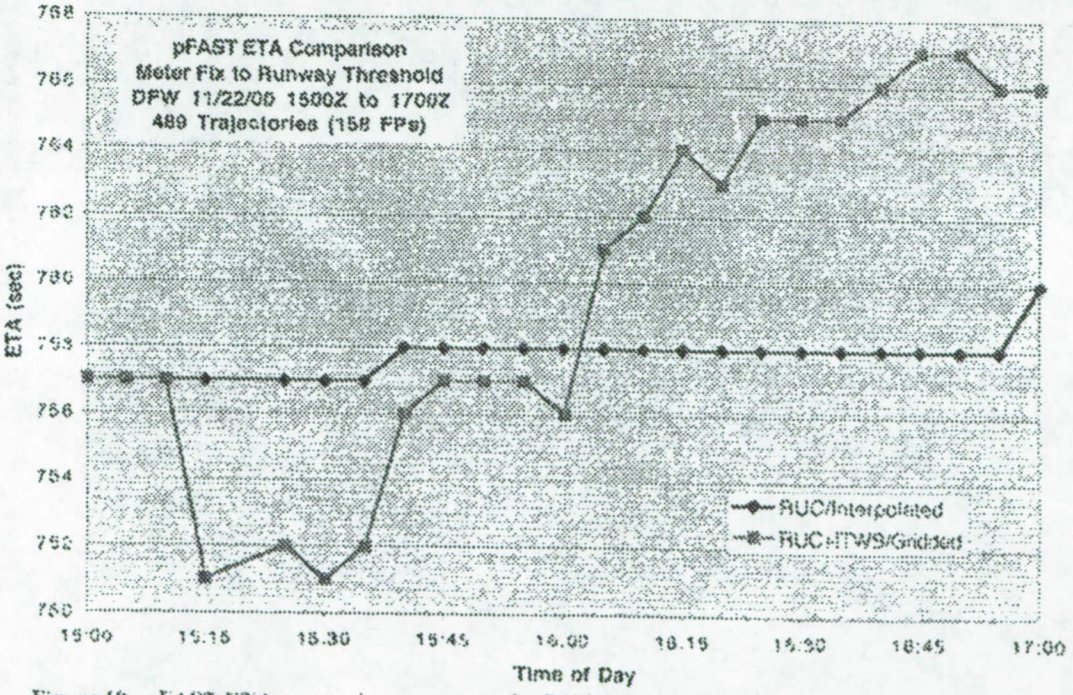

Figure 18. pFAST ETA comparison summary for RUC/Interpolated vs. RUC+ITWS/Gridded winds.

Table 2. Example performance measurement results (preliminary).

Process	Total Data Points, XxYxZ	CPU Use, seconds	Memory Use, Mbyte	Network Load, Mbyte		
				At the Hour	At the Half Hour	At 5 Minutes
WDPD*	30 x 16 x 55	11.73	24	30**		
	15 x 8 x 55	4.86	18	7**		
	9 x 5 x 55	3.37	18	2**		
WxServer*	30 x 30 x 50	10.98***	84***	0.7***	0.5***	0.2***
	24 x 24 x 50					
	70 x 70 x 70					
WxClient****		8.85***	34***			

*270 MHz Ultra 5 – 128 MB RAM
** Assumes 5 TS, 1 PFS, 2 PGUIs = Eight Using Applications
*** Total for all grids
****60 MHz SparcStation 20 – 64 MB RAM

SUMMARY

This paper described the design and implementation of the new CTAS weather distribution system. The approach relies on two key concepts. The first concept is the use of multiple nested wind grids for the TRACON, ARTCC and Adjacent Center airspaces to replace the single low-resolution weather grid currently used. This new method allows the use of higher spatial and temporal resolution products such as IWS Terminal Winds and improved RUC winds. The spatial resolution and update rate for each grid is tailored to the weather sources and user requirements.

The second concept is the use of nearest-neighbor data retrieval to replace the interpolation method currently used. Previous analysis showed that the memory requirements of the present method prevent its extension to higher resolution weather grids. The nearest-neighbor method was introduced in order to reduce the memory requirements to feasible levels. As an added bonus, the nearest-neighbor method also yields a substantial improvement in weather data access speed. Previous work showing that use of the nearest-neighbor method should not substantially degrade trajectory accuracy for the grid resolutions under consideration was confirmed in the present study.

The software design employs the concept of a single Weather Server process that provides weather data to multiple Weather Client processes. There is one Weather Server for a given CTAS site installation and one Weather Client for each workstation running one or more TS or PGUI processes. A reliable multi-cast protocol is used for transmitting the weather data from the Weather Server to the Weather Clients. The weather grids are divided up into subunits (called minor grids) and compressed for transmission using the GRIB format data compression technique. Each minor grid has a primary weather data source and optional secondary weather data sources. In the event that the primary weather data source is not available, the minor grid can continue to be generated using the secondary weather data sources.

The software is divided into ITWS Connection, Weather Server, Communications Protocol, Weather Client, Weather Library and Weather User modules. The ITWS Connection module inputs ITWS Terminal Winds and converts it to GRIB format. The Weather Server merges the RUC data (from the existing WDAD process) and the ITWS data to generate the nested grids information. The Communications Protocol module performs the reliable multi-cast of the nested grid data to the Weather Client processes. The Weather Client module receives the nested grid data and makes it

available to user processes via shared memory. The Weather Library provides the application program interface (API) for the user processes which selects the appropriate nested grid for data retrieval in a transparent manner. The Weather User module represents changes to the application processes where needed to accommodate the new method.

The testing procedures were also described. Results were provided including regression testing, performance measurement and software metrics. Finally, future work was described.

REFERENCES

1. Jardin, M.R., and H. Erzberger, "Atmospheric Data Acquisition and Interpolation for Enhanced Trajectory-Prediction Accuracy in the Center-TRACON Automation System", AIAA 96-0271, 34th Aerospace Sciences Meeting & Exhibit, Reno, NV, January 15-18, 1996.

2. Snyder, J.P., "Map Projections – A Working Manual", U.S. Geological Survey Professional Paper 1395, United Stated Government Printing Office, Washington, DC, 1987.

3. Dey, C.H., "The WMO Format for the Storage of Weather Product Information and the Exchange of Weather Product Messages in Gridded Binary Format", Office Note 388, U.S. Department of Commerce, NOAA, NWS, NCEP, 1996.

4. Lloyd, R.T., "CTAS Weather Processing – Data Acquisition, Conversion and Use", M.I.T. Lincoln Laboratory Project Memorandum 92PM-AATT-0004, Lexington, MA, 29 September 1999.

5. Kim, S. and R.E. Cole, "A Study of Time-To-Fly Estimates for RUC and ITWS Winds", M.I.T. Lincoln Laboratory Project Memorandum 95PM-Wx-0062, Lexington, MA, 21 January 2000.

6. Vandevenne, H.F. and J.R. Murphy, "Effect of Different Weather Formats on the Performance of Trajectory Synthesis in CTAS – A Case Study", M.I.T. Lincoln Laboratory Project Memorandum 92PM-AATT-0003, Lexington, MA, 14 March 2000.

7. Vandevenne, H.F. and J.R. Murphy, "Study of Trajectory Synthesis Efficiency with Respect to Weather Inputs, M.I.T. Lincoln Laboratory Project Memorandum 92PM-AATT-0005, Lexington, MA, 9 February 2000.

8. Vandevenne, H.F., "Grid Size of Weather Files vs. TS Performance", Annotated briefing, M.I.T. Lincoln Laboratory, 18 July 2000.

AIAA Guidance, Navigation & Control Conference, Montreal, Quebec, August 6-9, 2001 AIAA-2001-4362

EVALUATION OF ETA MODEL FORECASTS AS A BACKUP WEATHER SOURCE FOR CTAS*

Herman F. Vandevenne, Richard T. Lloyd and Richard A. Hogaboom

MIT Lincoln Laboratory
244 Wood Street
Lexington, MA 02420-9108

ABSTRACT

Knowledge of present and future winds and temperature is important for air traffic operations in general, but is crucial for Decision Support Tools (DSTs) that rely heavily on accurately predicting trajectories of aircraft. One such tool is the Center-TRACON Automation System (CTAS) developed by NASA Ames Research Center.

The Rapid Update Cycle (RUC) system is presently the principal source of weather information for CTAS.[1,2] RUC provides weather updates on an hourly basis on a nationwide grid with horizontal resolution of 40 km and vertical resolution of 25 mb in pressure.[3] However, a recent study of RUC data availability showed that the NWS and NOAA servers are subject to frequent service interruptions. Over a 210 day period (4/19/00-11/11/00), the availability of two NOAA and one NWS RUC server was monitored automatically. It was found that 60 days (29%) had periods of one hour or more where at least one server was out, with the longest outage lasting 13 hours on 9/21/00. In addition, there were 9 days (4%) for which all three servers were simultaneously unavailable, with the longest outage lasting 6 hours on 5/7/00. Moreover, even longer outages have been experienced with the RUC servers over the past several years.

RUC forecasts are provided for up to 12 hours, but these are not currently used in CTAS as back up sources (except that the 1 or 2 hour forecasts are used for the current winds to compensate for transmission delays in obtaining the RUC data). Since RUC outages have been experienced for longer than 12 hours, it is therefore necessary to back RUC up with another weather source providing long-range forecasts.

This paper examines the use of the Eta model forecasts as a back-up weather source for CTAS. A specific output of the Eta 32 km model, namely Grid 104, was selected for evaluation because its horizontal and vertical resolution, spatial extent and output parameters match most closely those of RUC.[4] While RUC forecasts for a maximum of 12 hours into the future, Eta does so for up to 60 hours. In the event that a RUC outage would occur, Eta data could be substituted. If Eta data also became unavailable, the last issued forecasts could allow CTAS to continue to function properly for up to 60 hours.

The approach used for evaluating the suitability of the Eta model and RUC forecasts was to compare them with the RUC analysis output or 0 hour forecast file, at the forecast time. Not surprisingly, it was found that the RUC model forecasts had lower wind magnitude errors out to 12 hours (the limit of the RUC forecasts) than the Eta model had. However, the wind magnitude error for the Eta model grew only from 9 ft/s at 12 hours (comparable with RUC) to 11 ft/s at 48 hours. We therefore conclude that RUC forecasts should be used for outages up to 12 hours and Eta model forecasts should be used for outages up to 60 hours.

METHODOLOGY

The comparison of RUC and Eta data was done for a typical ARTCC (Air Route Traffic Control Center), in this case the DFW Center (ZFW) airspace. In the vertical dimension, three altitude layers were examined: 7,500', 18,400' and 30,000'. The time period over which the comparison was made was a period of ten days, starting 9 June 2000.

Since the RUC and Eta model data is provided on different projection systems (Lambert Conformal vs. Polar Stereographic) and on different grids, the Eta data was first transformed onto the RUC grid using the following procedure.[5] Each RUC grid point was transformed into the Eta grid system and the surrounding Eta grid points determined. A linear interpolation was then done using the Eta weather products at the eight corners of the cube surrounding the RUC grid point. The resulting interpolated Eta model value was then used for the corresponding RUC grid point.

*This work was performed for the Federal Aviation Administration under Air Force Contract No. F19628-00-C-0002.

Once the data were on the same grid, metrics were adopted for comparing weather data. For temperature, the metric was simply the scalar difference between the RUC analysis temperature and the RUC or Eta model forecast temperature. For winds, there are two possible metrics: vector difference and magnitude/direction difference as illustrated in Figure 1.

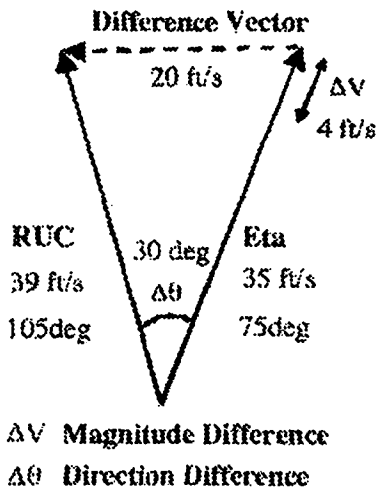

Figure 1. Metrics for comparing wind vectors.

The vector difference (shown by the dashed line in Figure 1) is often used as a wind error metric. However, this metric provides limited insight, since the orientation of the error vector varies greatly and the magnitude of the error vector is strongly influenced by the difference in direction between the two original vectors to be compared. For this study we chose the magnitude/direction metric ($\Delta V, \Delta\Theta$) as being more physically meaningful.

Finally, a selection must be made of the weather files to be compared. RUC outputs a set of forecasts on an hourly cycle, while Eta outputs a set of forecasts on a six hour cycle and while the longest RUC forecast is 12 hours, Eta forecasts as far as 60 hours into the future. Figure 2 shows the RUC and Eta cycles and the set of files generated at each update. The first file in each set is called the "analysis file" (abbreviated anl or r0 if for RUC and e0 if for Eta). This file represents the best knowledge of the weather at the analysis time after all the new measurements gathered during the previous cycle have been incorporated. The files following that are forecasts: a six-hour RUC forecast would be labeled r06 or simply r6 for RUC and e6 for Eta. A comparison between the RUC analysis file and an Eta 18-hour forecast would then be denoted by r0/e18 etc.

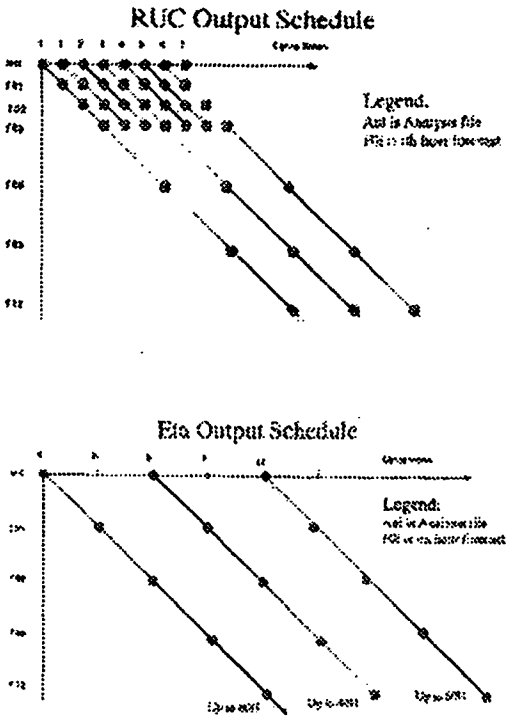

Figure 2. RUC and Eta model output schedules.

The comparisons that were made are first between the RUC and Eta analysis files, denoted r0/e0, for all sets output at the times corresponding to the Eta cycles (i.e. every 6 hours) for a period of 10 days. The 40 comparisons thus made were statistically analyzed separately and then combined for a global result in the end. Similarly, comparisons between RUC forecasts and a reference weather file, and Eta forecasts and the same reference file were made for the same 40 cycle times separately first and then the results were combined for a global result. The reference files were chosen to be the RUC analysis files that were generated at the time to which the forecasts were projected, i.e. a 12-hour forecast was compared with the r0 file generated 12 hours later. Although this seems to disadvantage Eta in a data quality comparison, it seemed justified by the fact that Eta is only a backup system to RUC and that RUC is the system normally in use. Finding an independent weather source (such as MDCRS readings etc) other than RUC or Eta to play the role of ground truth could be done but was beyond the scope of this study.

RESULTS

Figure 3 presents the results for the r0/e0, so-called "direct" data comparison. The x-axis is the index to the 40 Eta cycle updates and therefore the number of independent comparisons. The windows represent the results for comparison of temperature, wind strength and wind direction, averaged over the ZFW area and over the three selected altitudes.

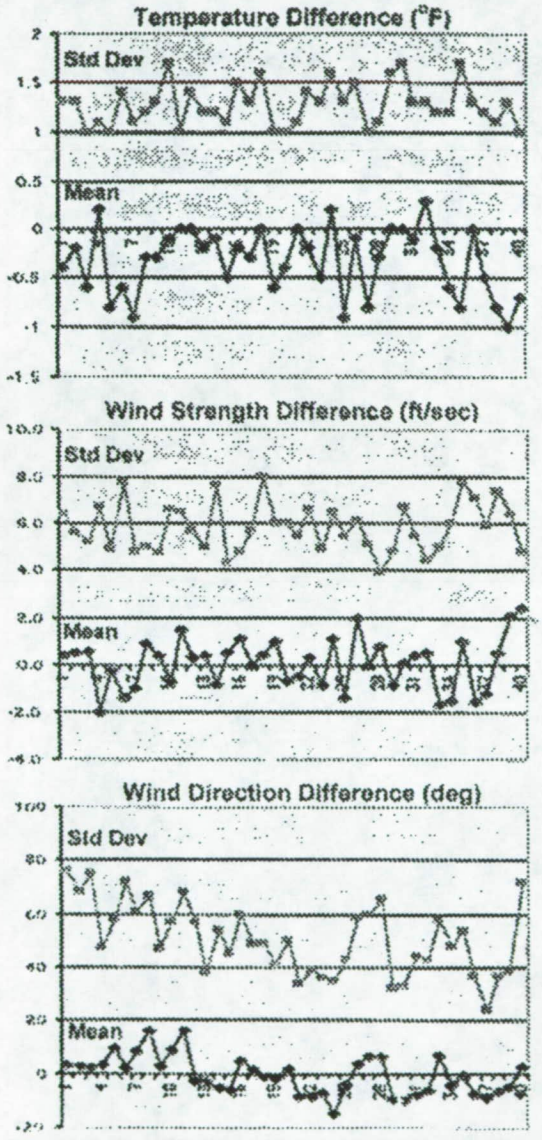

Figure 3. Direct RUC vs. Eta model comparison of temperature, wind strength and wind direction for 40 hourly updates (analysis cycles).

In order to better visualize what these results summarize we show in Figure 4 an overlay of the RUC and Eta wind fields at the altitude of 30,000ft (300mbar). One observes immediately that both systems present the same weather pattern, and that the wind strengths match well, but that there is a very discernible difference in wind direction. This seems to hold true for all comparisons made, whether among analysis files or forecasts. Wind strength varied from 5 ft/s in some parts to 45 ft/s in other parts, yet the average differences, returning to Figure 3, hover around zero, with a StD (standard deviation) of around 6ft/s for wind strength difference and a StD of about 1.25 degrees F in temperature difference. This seems to indicate that the differences between the RUC or Eta analysis files are small and that we are justified in taking either one (we choose RUC) as our reference when evaluating the RUC and Eta forecasts.

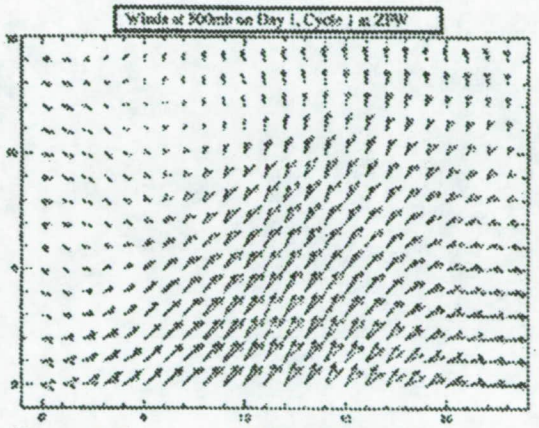

Figure 4. Overlay of RUC and Eta wind vectors.

The next comparisons are between RUC forecasts and the RUC analysis data, generated at the time to which is forecasted; and similarly between Eta forecasts and RUC analysis weather that materialized later at the appropriate time.

Figure 5 shows the results for the 12hour forecasts of the temperature: one curve for r0/r12 and another for r0/e12, with as x-axis the 40 Eta update times. In the first window we show the mean difference and in the second the StD of the temperature difference. In Figures 6 and 7 we show similar results for the wind strength and wind direction differences. A close look reveals that peaks and valleys in both curves match quite closely. This means that there were weather changes not predicted by either system, and that they erred in the similar ways.

Figure 8 shows means and StD for the comparison r0/e48, for the 48hour Eta forecast (no curve for RUC can be shown since RUC only forecasts for up to 12 hours).

Figures 9,10 and 11 summarize all these results averaged over all 40 updates. These new figures contain the mean and StD of differences r0/ri for 1,2,3,6 and 12 hours and r0/ej for 0,6,12,18...48 hours for the parameters temperature, wind strength and wind direction.

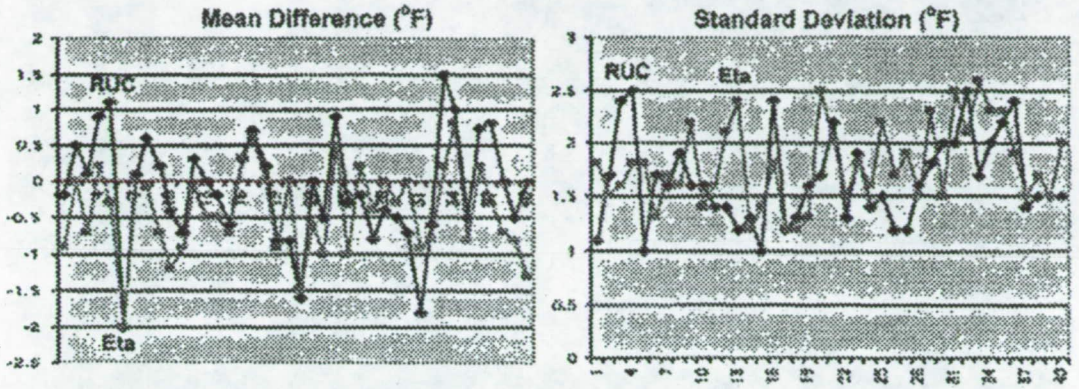

Figure 5: Comparison of 12-hour temperature predictions for RUC and Eta forecasts vs. RUC analysis data for 40 analysis cycles.

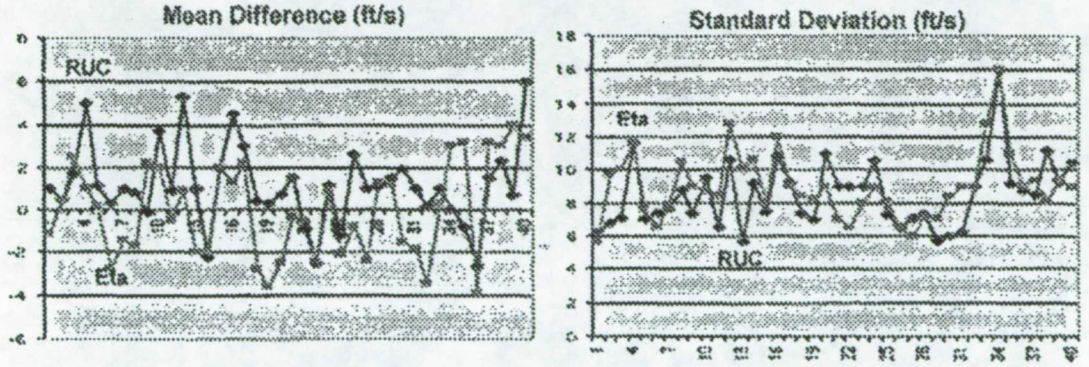

Figure 6: Comparison of 12-hour wind strength predictions for RUC and Eta forecasts vs. RUC analysis data for 40 analysis cycles.

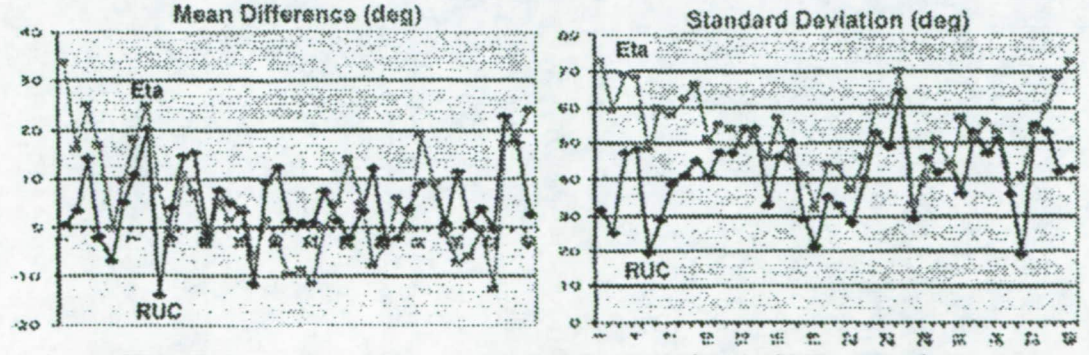

Figure 7: Comparison of 12-hour wind direction predictions for RUC and Eta forecasts vs. RUC analysis data for 40 analysis cycles.

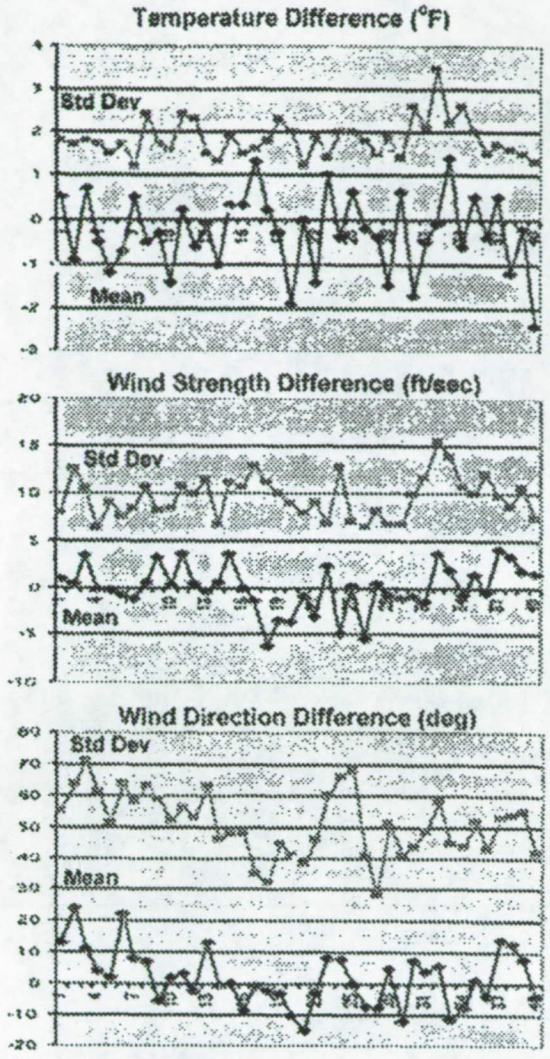

Figure 8: Evaluation of Data Quality of the 48-Hour Eta Forecasts

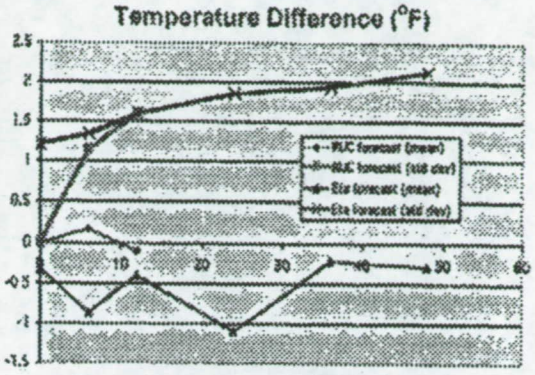

Figure 9: Summary of temperature differences for RUC and Eta forecasts vs. RUC analysis (°F).

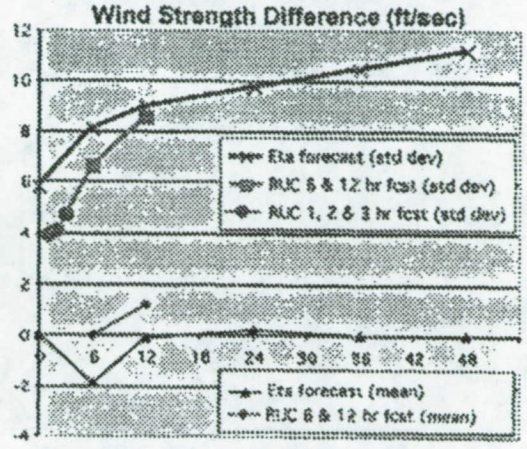

Figure 10: Summary of wind strength differences for RUC and Eta forecasts vs. RUC analysis (ft/s).

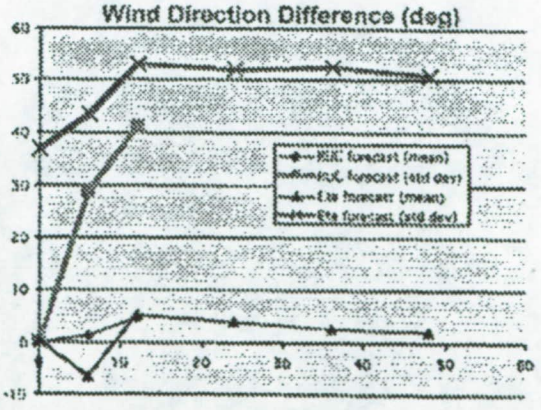

Figure 11. Summary of wind direction differences for RUC and Eta forecasts vs. RUC analysis (deg)

We observe that, based on the StD, one could state that RUC forecasts are marginally better than the comparable Eta forecasts for the same duration into the future up to 12 hours, at which point they are the same. But one can also see that 6 hour Eta forecasts are as good or better than 12hour RUC forecasts. The Eta forecasts maintain almost the same relative quality up to the maximum duration of 48hours tested (and by extension to the 60-hour duration available from Eta). Observe that the StD of temperature differences stays below 2 degree Fahrenheit for forecasts up to 48 hours, and StD of wind strength differences stay below 11ft/s (6.5knots) for even the maximum duration forecasts. The wind direction difference, although with mean about zero, has a StD of a steady 50 degrees. These and some more detailed observations form the basis for the proposed switching protocol when either RUC alone or both RUC and Eta data become unavailable

CONCLUSIONS

The purpose of this study was to verify the quality of the Eta forecasts and to propose a decision algorithm for switching from RUC to Eta in case of outages. First let us state that Eta forecast quality is not in doubt in view of the results presented in the previous paragraph. Next, the protocol for switching will be based on the presumed quality of the forecasts. For example, if access to both RUC and Eta is denied simultaneously, one should continue running CTAS with RUC forecasts as long as possible (from 9 to 12 hours, depending when in the RUC cycle the outage occurred). If RUC is interrupted, but not Eta, then one should switch at the next Eta update that would be at most 6 hours after the interrupt. There are some additional considerations: CTAS expects a new RUC weather file to be made available every hour. When interrupts occur it may be necessary to create hourly files by interpolating from two adjacent forecast files. This would be true for RUC after the third hour and is always true for Eta forecasts as is clear from Figure 2 showing the RUC and Eta output products. The complete algorithm depends on the exact time of start of the outage compared to the underlying RUC and Eta cycle time, but it is based on the presumed quality of RUC and Eta forecasts at any given time.

REFERENCES

1. Jardin, M.R., and H. Erzberger. "Atmospheric Data Acquisition and Interpolation for Enhanced Trajectory-Prediction Accuracy in the Center-TRACON Automation System", AIAA 96-0271, 34th Aerospace Sciences Meeting & Exhibit, Reno, NV, January 15-18, 1996.

2. Campbell, S et al, "The Design and Implementation of the New Center/TRACON Automation System (CTAS) Weather Distribution System", AIAA GN&C Conference (AIAA-2001-4361), Montreal Quebec, Canada, 6-9 August 2001.

3. Benjamin, S. "Present and Future of the Rapid Update Cycle", Slide Presentation, CWSO National Conference, http://maps.fsl.noaa.gov/CWSU/ruc_feb00_C_WSU_files/v3_slide0001.htm, 9 February 2000.

4. "ETA Model Characteristics: Background Information", COMET/COMAP Training Symposium on Numerical Weather Prediction http://www.comet.ucar.edu/nwplessons/etalesson2/characteristicsbackground.htm, December 1998

5. Snyder, J.P., "Map Projections – A Working Manual", U.S. Geological Survey Professional Paper 1395, United Stated Government Printing Office, Washington, DC, 1987.

AIAA Guidance, Navigation & Control Conference, Montreal, Quebec, August 6-9, 2001 AIAA-2001-4363

WEATHER IMPACTED ROUTES FOR THE FINAL APPROACH SPACING TOOL (FAST)*

James R. Murphy
NASA Ames Research Center
Moffett Field, CA 94035-1000

Steven D. Campbell
MIT Lincoln Laboratory
Lexington, MA 02420-9185

ABSTRACT

This paper addresses the issue of developing weather-impacted routes for the Final Approach Spacing Tool (FAST). FAST relies on adaptation data that includes nominal terminal area routes and degrees of freedom to generate optimum landing sequences and runway assignments. However, during adverse weather some adapted routes may become unavailable due to the presence of hazardous weather. If FAST continues to generate trajectories using these routes, its schedule will not be accurate during the adverse weather. The objective of the study was to determine methods for incorporating severe weather products and weather-impacted route data into FAST.

INTRODUCTION

Increasing air traffic demand in the face of limited runways and airspace has made improving the efficiency of the nation's air traffic system one of the Federal Aviation Administration's (FAA) top priorities [1]. New decision support tools are being developed to assist Air Traffic Controllers and Traffic Managers in achieving these efficiency gains while maintaining safety.

One such tool is the Final Approach Spacing Tool (FAST), which is an element of the Center/TRACON Automation System (CTAS) being developed by NASA Ames Research Center. There are two versions of FAST: Passive FAST (pFAST) and Active FAST (aFAST). Passive FAST provides controllers with runway assignments for delay balancing and landing sequence numbers to optimize the landing order. Active FAST additionally provides controllers with heading, speed and altitude advisories to achieve these optimal sequences [2]. Passive FAST has been operationally tested at the Dallas/Ft. Worth International Airport (DFW) and has demonstrated up to a 13% airport throughput increase [3].

Active FAST, which is currently under development, is expected to offer additional increases in airport capacity. For the remainder of this paper, "FAST" will be used generically to refer to both systems.

FAST relies on adaptation data that includes nominal TRACON routes and degrees of freedom to generate optimum landing sequences and runway assignments. However, during adverse weather, some adapted routes may become unavailable due to the presence of hazardous weather. If FAST continues to generate trajectories using these routes, its schedule will not be accurate during the bad weather. The objective of the study was to determine methods for incorporating severe weather products and weather-impacted route data into FAST.

Air Traffic Control (ATC) decision support tools have many sources of weather data. Table 1 outlines some of these data sources and the information supplied.

Since FAST operates in the terminal area, the key source for weather information for this study was the Integrated Terminal Weather System (ITWS). ITWS was developed under FAA support by MIT Lincoln Laboratory (MIT/LL) and is currently being implemented for deployment at 35 locations across the continental United States. ITWS integrates weather data from a variety of sources, including Terminal Doppler Weather Radar (TDWR), Next Generation Weather Radar (NEXRAD), Airport Surveillance Radar (ASR-9), Rapid Update Cycle (RUC) model and surface sensor information. It generates a variety of weather display products for Traffic Management Coordinators (TMCs), controllers, airlines and other users. These products have been operationally tested at several sites including Dallas/Ft. Worth, Atlanta, Denver, Los Angeles and Miami, and have demonstrated significant operational benefits by allowing FAA personnel to improve their ability to manage traffic during bad weather [4].

*This work was performed for the National Aeronautics and Space Administration under Air Force Contract No. F19628-00-C-0002

This study focuses on the integration and use of various ITWS weather products in FAST. It also addresses the display of an integrated system to the end-user. In this case, identified users of the system would be the Traffic Managers at either the TRACON or the Air Route Traffic Control Center (ARTCC). The italicized ITWS products from Table 1 will be emphasized in this paper.

APPROACH

A study by Rhoda and Pawlak [5] revealed pilot behavior in choosing to penetrate or deviate around hazardous weather depends on several factors. Conventional wisdom is that pilots deviate around Video Integrator and Processor (VIP) level 3 or higher weather. However, in evaluating 1,952 encounters with weather at DFW, it was found that a significant proportion of the encounters (>10%) resulted in penetration of VIP level 3, 4 or 5 weather. It was also found that the closer an aircraft was to the airport, the more likely it would penetrate VIP level 3 or higher weather. A neural net classifier was developed to create Probability of Deviation (PODEV) maps giving the probability of a typical aircraft penetrating a given region of airspace.

With the development of the Terminal Convective Weather Forecast (TCWF), Wolfson, et al. [6] showed a significant increase in skill could be achieved in forecasting the location of VIP level 3 and higher weather for line storms up to one hour in the future. TCWF has been running experimentally at DFW since the summer of 1999. This product is generated in conjunction with the ITWS testbed operated by MIT/LL at DFW. The TCWF has been proposed as a Pre-Planned Product Improvement (P3I) to ITWS

By applying the TCWF algorithm to the forecasted weather products in ITWS, up to one hour of predicted VIP level 3 or higher weather probability maps at can be generated at 10-minute intervals. The neural net classifier used to create the PODEV maps can then be applied to the TCWF forecasts to generate Forecasted PODEV (FPODEV) maps. The PODEV and FPODEV maps are overlayed with the FAST adapted aircraft routes to generate the probability of weather impact for each route segment from the present to one hour in the future in 10-minute intervals.

An experimental weather and traffic Graphical User Interface (GUI), referred to as the Offline Traffic and Weather Display (OTWD), was developed to facilitate use of the weather-impacted route information. The aircraft, weather products and FAST adaptation routing information are displayed on OTWD. The route segments are color-coded to reflect the probability of weather impact at a particular time. A time slider was created to allow the display of the weather forecast and show the weather impact on aircraft route segments. The user can also select any route segment to determine the onset and duration of hazardous weather

Table 1. Weather Input Data Types

Data Source	Weather Types	Coverage
ITWS	Winds *Reflectivity* *Probability of Deviation* *Convective Forecasts* Gust Fronts Microbursts	Terminal
NOAA	HUC Winds	CONUS
WSI	Reflectivity Lightning Jet Stream	CONUS
ETMS	Reflectivity Lightning Jet Stream	CONUS

An experimental version of the Airport Configuration Manager (ACM) was developed to facilitate TMC use of the weather-impacted routes information. ACM shows the impact upon the schedule of weather-impacted routes by predicting future airport configuration changes for each adapted airport. This capability alerts the TMC to weather changes in the terminal area that could affect the operation of the airport up to one hour in advance. It also suggests an appropriate future airport configuration based on predicted weather data.

Figure 1 demonstrates possible uses of ITWS weather products in FAST, from acquisition and parsing of the data to predicting impact on aircraft and terminal area operations. ITWS created precipitation maps can be converted into PODEV maps, which then can be used to identify routes and aircraft that are impacted. Delays caused by the hazardous weather can then be considered for ATC decisions such as airspace closures and re-routes. In addition, severe weather forecasts can be used to predict impacted aircraft and routes up to an hour in the future, as well as help determine appropriate future airport configurations. Schedule changes based on the weather forecasts can be utilized by TMCs to facilitate further ATC decisions.

RESULTS

OFF-LINE TRAFFIC AND WEATHER DISPLAY

The identification and display of weather-impacted routes has been developed and demonstrated both on OTWD and on the FAST Planview Display GUI (PGUI), see Figure 2 [3]. Initial evaluation was completed using feedback from MIT/LL in-house controllers resulting in the current GUI design. The development of OTWD promotes the operational concept evaluation by allowing for the rapid prototyping of display modifications. However, the current functionality has only been implemented as a playback system. Additional work is needed to implement these functions into a real-time tool and to evaluate the initial operational concept.

OTWD displays aircraft and FAST terminal area routing, as well as the various weather products. The tool reads FAST generated playback files and synchronizes its internal clock to the data file time. As the weather products are read into the system at the appropriate time, OTWD analyzes the routing structure and weather data to determine the time, duration and severity of impaction on each route segment. These route segments make up the nominal paths aircraft use to fly from the meter fix to each runway threshold.

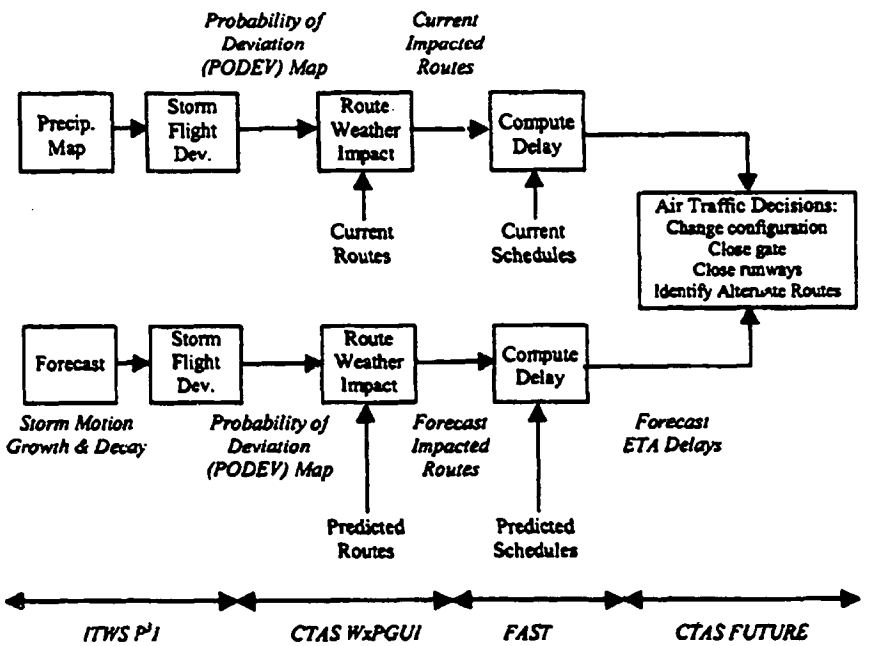

Figure 1. Weather Impacted Routes Processing

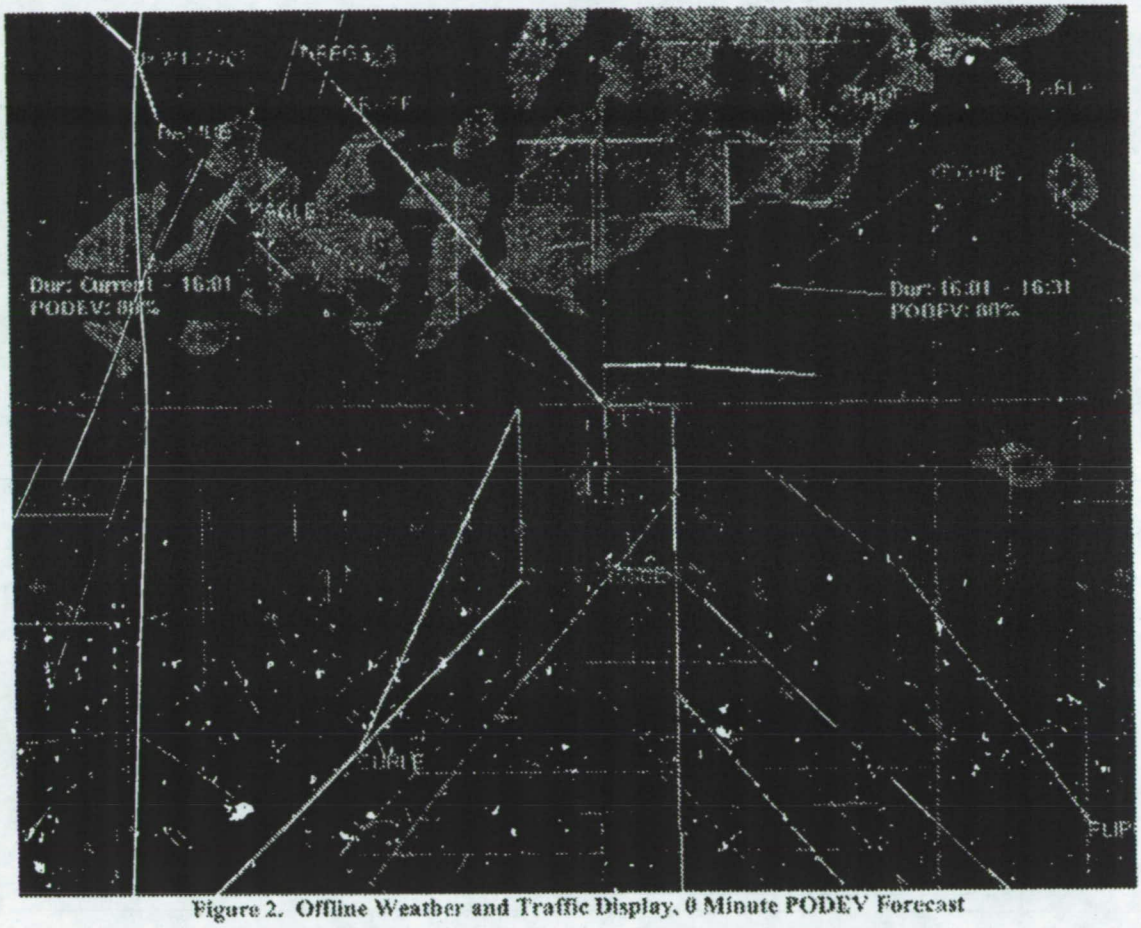

Figure 2. Offline Weather and Traffic Display, 0 Minute PODEV Forecast

Figure 2 shows an example of a PODEV map for the DFW TRACON at 15:53 GMT. The contours range in intensity level from 20% - 100% shown by increasing intensity of shaded areas. Notice the selected route segment starting at the fix "BAMBE". It has a predicted duration of impact from "Current" to 16:01, with a Probability of Deviation of 80%. This means that from 15:53 GMT until 16:01 GMT, a typical flight could be expected to deviate from this route segment 80% of the time. The information displayed for the route segment beginning at "COVIE" shows that starting at 16:01 GMT through 16:31 GMT it will have a PODEV of 80%, though this route segment is not currently impacted by severe weather.

The future weather forecast used to predict the airspace PODEVs can also be displayed on OTWD using the slider bar, as seen in Figure 3. This figure shows the FPODEV at 16:14 GMT, with the traffic data from 15:54 GMT allowing the user to visualize how airspace problems might develop over time. Though not currently part of OTWD, the display of the predicted aircraft positions at the "look ahead" time has been demonstrated and could be beneficial.

Since PODEV maps are derived from 6 VIP level and TCWF data, OTWD is also capable of displaying these data. Figure 4 shows the TCWF for the DFW TRACON. The shaded areas correspond to predicted areas of VIP level 3 or higher weather activity twenty minutes from the current time.

One area of research is to extend TCWF to other areas of the country. TCWF has been tested and works well with Convective storms. However, some areas of the country tend to have more localized thunderstorm cells. Forecasting these types of weather cells is still being researched. Hence, the PODEV and terminal area forecasts for some areas of the country may not be mature enough to predict impacted routes accurately.

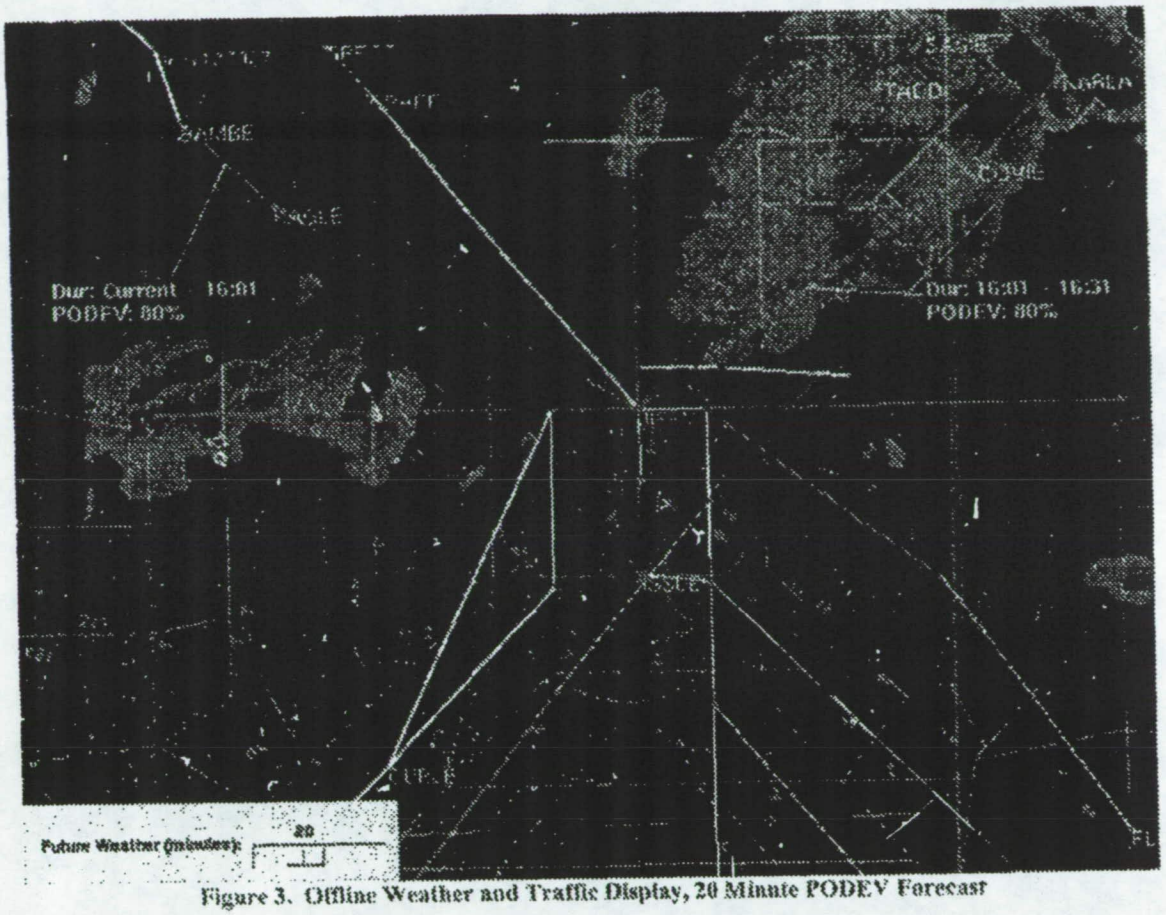

Figure 3. Offline Weather and Traffic Display, 20 Minute PODEV Forecast

AIRPORT CONFIGURATION MANAGER (ACM)

Using the identified impacted routes, an algorithm to help identify the most appropriate airport configuration for the current and predicted weather scenarios was developed. FAST uses the concept of airport configuration not only to identify the actual runways being used for landing, but also to define the different procedures used to control aircraft through the TRACON. Taking the current and forecasted weather information into account, ACM is designed to identify appropriate future configuration changes and determine their time and duration.

The most important aspect for determining the appropriate airport configuration is the condition at the runway threshold. This includes the strength and direction of the wind. However, the ITWS weather data does not currently include terminal area wind predictions. Therefore, the search for future airport configurations is limited to the configurations adapted for FAST with runways landing in the same direction as the current configuration. The algorithm necessarily assumes that the wind direction will remain relatively constant.

First, the current airport configuration is identified. The adapted configurations are searched for others that contain any of the runways used in the current configuration. This forms our basic list of choices for future configuration changes. All runways used by any configuration in this list comprise our selection of possible future runways.

Next, the time when the current runways in use will change needs to be determined. To do so, the adapted final approach segment for each runway from our possible future runways list is examined. For each 10-minute forecast interval over the next hour, it is determined whether each runway is impacted or not. Once the list of forecasted non-impacted runways differs from the runways utilized in the current airport configuration, a future configuration change is identified.

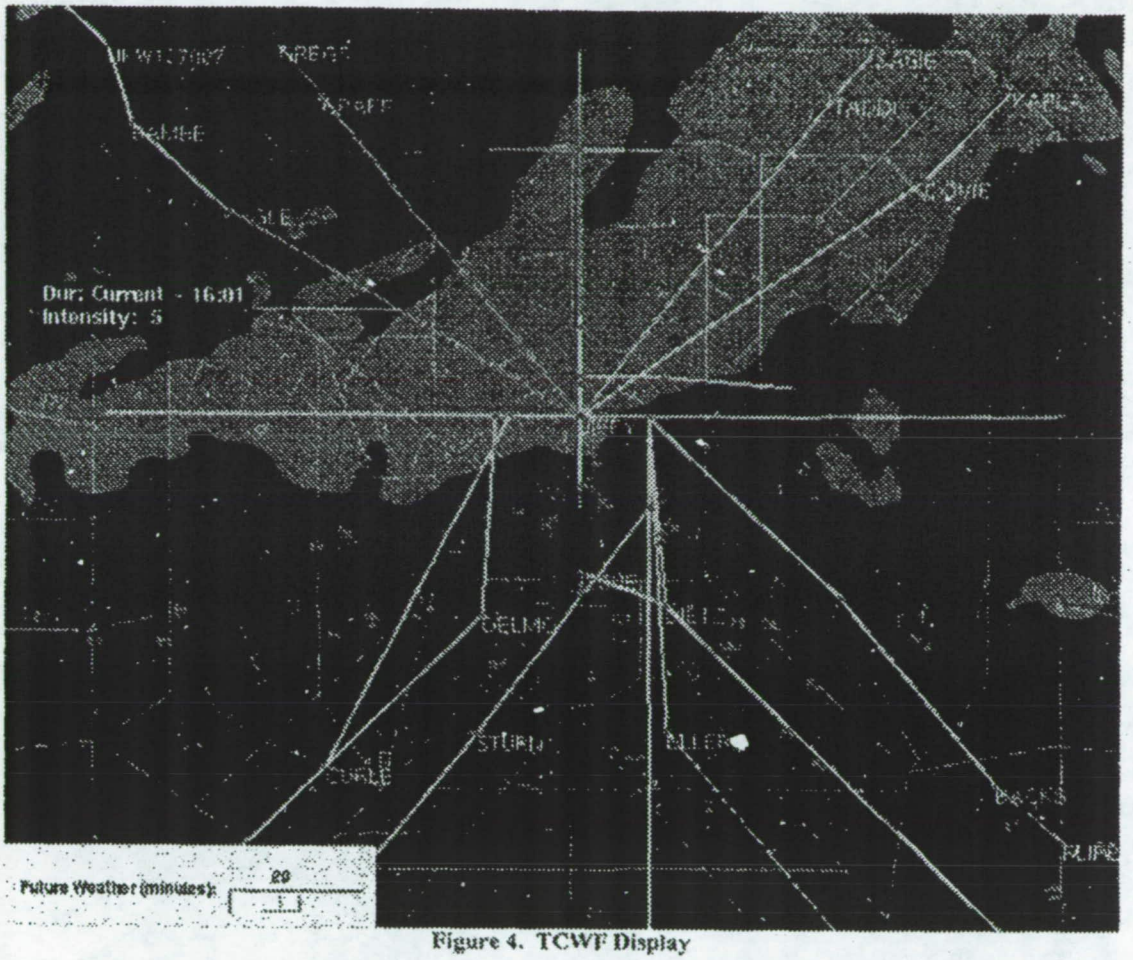

Figure 4. TCWF Display

In practice the changing of airport configurations can be difficult to perform. An arrival aircraft is assigned to a runway upon entering the TRACON. An aircraft can take up to 20 minutes to fly through the TRACON under normal conditions, so changing the airport configuration too often would not be practical. Therefore, future configuration changes are not automated in the system. Instead, a possible change is indicated, and the Traffic Managers can use the information to determine the best course of action.

NEXT STEPS

Immediate next steps include implementing a real-time version of the OTWD at DFW. One component of the implementation would be to create the infrastructure and interface to the real-time ITWS data at DFW. Integration work has already begun for the ingestion of ITWS wind data to augment the RUC forecast data. The inclusion of additional weather products would be straightforward to implement. PODEV maps are currently not included in the operational ITWS system and either another source of these maps would need to be created, or the software to create the PODEV maps would need to be integrated into FAST. Once the impacted route software is in place, the value of AMC could be evaluated using operational TMCs at the DFW TRACON

The use of ITWS wind forecasts to help manage the airport configuration could also be investigated. As stated above, ACM uses the runways available for aircraft landing in the current direction to determine future runway usage. However, it would be more useful to include predictions of changes in the wind direction that could impact the direction aircraft can land.

As an extension to ACM, the identification and notification of impacted meter fixes to FAST could be added. This would allow the Traffic Managers to close meter fixes up to an hour in advance, giving aircraft in the En Route airspace the opportunity to re-route.

OTWD could also be expanded from the terminal area to the en route area. With the incorporation of other weather data sources such as turbulence maps and en route weather forecasts, the same tools could be applied to help predict airspace closure and aircraft intent outside the terminal airspace. This could improve predictions necessary for conflict detection and airspace overloading.

Once the prediction of weather-impacted airspace and aircraft routes is complete, a logical next step would be to identify specific aircraft with impacted trajectories. This activity would tie together the predicted configurations and weather-impacted routes development as it pertains to the runway assignment of aircraft in the terminal area. Aircraft could be color coded by weather-impact to alert the controllers to the need to utilize alternate routes.

An initial implementation could allow TMCs to graphically modify weather-impacted routes to create alternate nominal route segments for FAST to use for determining the path of flight in the terminal area. This would side step the procedural complexity of managing arrival and departure airspace.

Based on route-building algorithms and heuristics investigated by Krozel et. al.[7], the indirect effects of severe weather on the closure of airspace due to thunderstorm downstream turbulence and severe weather in missed approach areas could be researched [8]. From these analyses, FAST could ultimately develop the alternate routes automatically.

CONCLUSIONS

A method for connecting current and forecasted hazardous weather with FAST has been developed. The method utilizes previous work for creating probabilistic maps of pilot behavior in penetrating or deviating around hazardous weather. This work has been extended by applying a new method for forecasting up to one hour in advance the probability of significant weather in a given region. These two methods are combined to produce forecast maps of probability of deviation for up to one hour in advance. These probability maps are applied to FAST adaptation data to determine current and forecasted weather-impacted routes. These weather-impacted routes are used in an interactive display that could be used to provide guidance to Traffic Managers in choosing airport configurations in the presence of hazardous weather.

REFERENCES

1. Federal Aviation Administration, "FAA Strategic Plan." FAA Document A-1, January 2001.

2. Robinson III, J. E., D.R. Isaacson, "A Concurrent Sequencing and Deconfliction Algorithm for Terminal Area Air Traffic Control", AIAA Guidance, Navigation, and Control Conference, Denver, CO, August, 2000.

3. Davis, T. J., D. R. Isaacson, J. E. Robinson III, W. den Braven, K. K. Lee, and B. Sanford, "Operational Field Test Results of the Passive Final Approach Spacing Tool", IFAC 8th Symposium on Transportation Systems, Chania, Greece, June, 1997.

4. MIT Lincoln Laboratory, "Integrated Terminal Weather System (ITWS) Algorithm Description." Rev B. DOT/FAA/ND-95/11. LL Project Report No. ATC-255, 1998.

5. Rhoda, D. A. and Pawlak, M. L., "An Assessment of Thunderstorm Penetrations and Deviations by Commercial Aircraft in the Terminal Area." Project Report NASA/A-2, June 1999

6. Wolfson, Marilyn M., B.E. Forman, R.G. Hallowell, and M.P. Moore, "The Growth and Decay Storm Tracker," AMS 8th Conference on Aviation, Range, and Aerospace Meteorology, Dallas, TX, January 1999.

7. Krozel, Jimmy, T. Weidner, and G. Hunter, "Terminal Area Guidance Incorporating Heavy Weather," AIAA Guidance, Navigation, and Control Conference, New Orleans, LA, August 1997.

8. Lester, P.F., "Turbulence: a new perspective for pilots," Jeppesen Sanderson. Englewood, CO., 1993.

AIAA Guidance, Navigation & Control Conference, Montreal, Quebec, August 6-9, 2001 AIAA-2001-4360

USING SURFACE SURVEILLANCE TO HELP REDUCE TAXI DELAYS*

Jerry D. Welch and Steven R. Bussolari
MIT Lincoln Laboratory
244 Wood Street
Lexington, MA 02420-9108

Stephen C. Atkins
NASA Ames Research Center
Mail Stop 210-
Moffett Field, CA 94035-1000

ABSTRACT

Taxi delay is the largest of all aviation movement delays. However, taxi-out delays have not received attention equal to that focused on airborne delays because taxi-out delays often result from downstream problems. Also, until recently, there was no practical means of tracking surface movements. New surface surveillance technology will revolutionize surface management by providing data for planning, timing, and monitoring surface operations. This paper proposes a simple aid to help manage departure taxi queues and help exploit existing departure capacity, while avoiding the delays that result from saturated queues and unbalanced runways. The proposed decision aide will use archived surveillance data to quantify queuing behavior and model departure capacity, and it will use real-time surveillance to track capacity changes and monitor the state of the taxi queues.

INTRODUCTION

Taxi delay results in the largest direct operating cost to US air carriers of all delays. Fig. 1 shows that the average taxi-out delay in minutes per flight is approximately twice the airborne delay.[4] Although aircraft burn fuel roughly 5 times faster when airborne, crew and equipment costs make the spend-rate for taxiing aircraft about 2/3 that for airborne aircraft. Consequently, the cost of taxi-out delay exceeds that for airborne delay by about 1/3, totaling more than one billion dollars annually.

On average, taxi-out delay is 3 times larger than taxi-in delay. This situation suggests that surface aids will likely focus first on departures. Most taxi-out delays are associated with surface queuing processes that are visible to the surface surveillance system. In contrast, management of the taxi-in process would require access to airborne surveillance data. Therefore surface planning aids for arrivals would need to be integrated that focus on with ARTS terminal surveillance data or the Center-TRACON Automation System (CTAS).[11]

Surface aids departures could rely solely on surface surveillance and flight plan data, and could operate independently of ARTS and CTAS.

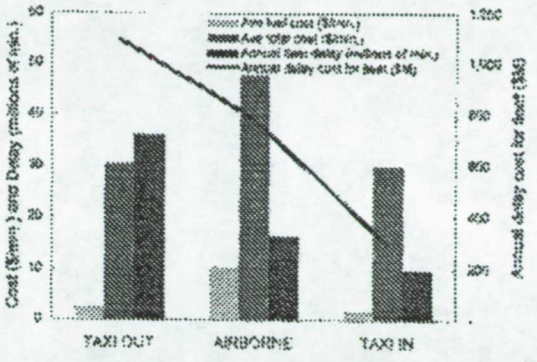

Figure 1. The cost of delay - U.S. airline fleet cost estimates for 1999 from (ATA, 2000).[4]

Efforts to reduce taxi-out delay have not been as vigorous as those focused on airborne delay. One likely reason is that taxi-out delay often stems from downstream problems such as capacity limitations in terminal and en route airspace as well as at the destination airport. Delay statistics show a 70% correlation between taxi-out delay for departures from DFW and airborne delay for flights to DFW.[5] This means that, on days when arrivals incur longer delays, departures at the same airport usually take more time to get off the ground. Consequently, there has been a sense that little can be done on the surface to reduce taxi-out cost.

The airport surface has been the only domestic aviation domain without an automatic means of digitally tracking and identifying aircraft. There has been no practical means of tracking surface movements to understand or control the taxi process, other than using reported push-back and wheels-off times.

High performance cooperative surface surveillance with aircraft identification has now been demonstrated.[20] This new technology will provide an opportunity to revolutionize air traffic management on the airport surface. There has been considerable research in the application of surface surveillance for enhancing

*This work was performed for the National Aeronautics and Space Administration under Air Force Contract No. F19628-00-C-0002.

surface safety.[8, 19] There has been some research in analyzing and understanding taxi delays.[12] An attempt was made to mitigate taxi delays by predicting queue lengths from schedule data.[13] However, no work has been reported on the use of surface surveillance to help enhance surface efficiency and reduce taxi delay. Surface surveillance will help improve the understanding of taxi delay mechanisms and will provide data for characterizing operational constraints, planning taxi operations, timing taxi clearances, and monitoring aircraft compliance.

The airport surface has characteristics that make it attractive for implementing automatic surveillance-based decision support tools. The surface is the only aviation domain in which most aircraft follow a relatively small number of rigidly constrained paths. This may make it feasible to use surveillance information alone to evolve a database that fully characterizes the geometry of the taxi paths. Generic adaptation algorithms could be designed to use this taxi path data to automatically adapt surface decision support tools to each new airport site. Automatic adaptation would significantly reduce the cost of implementing surface decision support tools relative to terminal and en route tools.

It appears feasible to develop operational software that uses surface surveillance information alone to track runway configuration changes in real time. By adding flight plan data to the surveillance data, automatic algorithms should be able to determine the current surface queuing delay status for each runway and aircraft. Other databases could relate taxi delay to demand and relate demand to the day of the week and the time of day. Operational algorithms could use these delay and demand databases in conjunction with measurements of departure queuing status to automatically predict the near-term departure throughput for each runway. With the addition of surveillance data for arriving aircraft, operational algorithms could also predict near-term arrival throughput.

A well-designed surface surveillance system can provide complete coverage of the important delay-controlling queues on the airport surface. Surveillance coverage of airport surface queues can provide the visibility into surface traffic flow that is essential to close the loop on suggested control actions. Surface surveillance provides a means of determining departure throughput performance. It allows unambiguous determination of the departure runway of each aircraft. By contrast, the ARTS surveillance system cannot reliably associate departures with runways because of its low-elevation coverage cut-off. This is particularly true when the departures turn immediately after takeoff.[3]

A simple taxi-out aid could use this surveillance information to predict the queue length for the next few minutes and advise optimum near-term target pushback rates. Such advisories would have the advantage that they would not require any manually generated information from controllers or aircraft operators. There would be no need for controllers to provide the current runway configuration and no need for pushback predictions from aircraft operators.

A SIMPLE TAXI-OUT AID

The most elementary taxi aid would display surface traffic with aircraft identities to all parties interested in surface traffic management. Distributing surveillance information with aircraft tags would likely benefit aircraft operators and controllers and would require little research other than finding means to manage the tags to avoid display clutter. The functionality envisioned in this paper goes further and provides decision support aids based on surface surveillance that would attempt to directly help controllers and air carriers work together to improve surface movement efficiency. This is done by providing a simple visualization of the queuing situation along with pushback advisories to optimize the taxi-out process and balance runways.

At airports with multiple runways, surveillance data alone presents a good tactical view of the airport, but sometimes paints a confusing picture of the strategic situation. Early in the taxi-out process, it is not always clear which runway each aircraft is heading for. Although the total number of taxiing aircraft may be apparent, it is difficult to see the queue lengths for individual runways. The surveillance display also does not provide information on predicted taxi times.

Experience with the Surface Movement Advisor (SMA) program at Atlanta's airport conclusions from current NASA-sponsored research on causes of departure delay and inputs from aircraft operators who routinely experience departure delays all suggest that an automation aid to help visualize and control taxi queue lengths for departing aircraft would benefit both controllers and aircraft operators.[5,13,16,12,14,1,2,3,9,8]

This paper proposes such a simple taxi-out aid to help with queue management. Figure 2 shows a notional operator interface for the aid when used at an airport with dual departure runways. The interface is presented merely as a concrete illustration of the proposed

information content. It has not been prototyped or tested for operator acceptability.

It displays the status of the taxiway system and depicts the queuing situation at multi-runway airports by associating each aircraft with a runway and depicting the airport in a simplified map format similar to those used for mapping subway lines. It shows the loading of the taxi paths as well as the runway queues, indicating the progress of all aircraft taxiing towards each departure runway.

The runway maps are divided nominally into Pushback, Taxi, Runway Crossing Queue, and Runway Queue sections. The taxi and Queue sections contain bins consisting of columns of boxes indicating numbers of aircraft. Each box corresponds to an aircraft. Clicking on a box displays a data tag with the aircraft ID. In the figure, runway 18C has two planes waiting to take off in its runway queue. Where a taxi path crosses an active runway, the graph is split by a crossing queue as shown for Runway 18R.

The Taxi section is shown with 10 bins. Typically, each taxi bin contains no more than one aircraft. Large airports use more taxi bins to keep the individual taxi bin occupancy to a single aircraft.

A Crossing Queue section has a single column showing the number of aircraft waiting to cross the active runway. When an aircraft first enters a runway crossing queue or a runway queue, its box changes color to indicate that it is now under local control. Finally, the graph for each runway ends in a Runway Queue section containing a single column showing the number of aircraft waiting at the runway entrance for takeoff clearances.

Each runway also has a pushback rate advisor. When the departure demand is low, there is no need to limit the pushback rate. Thus, although the display operates continually, its principal benefit occurs in departure push conditions. The height of the bar recommends a pushback rate for that runway. A positive bar indicates that the pushback rate can be increased, and a negative bar advises a reduced rate. In a departure push the goal is to zero the bars for both runways. Zeroing the bars optimizes the pushback rates and equalizes the taxi times for the last aircraft in each of the runway queues. (Note that the last aircraft in a runway queue may not be the most recent aircraft to push back for that runway.) The display also indicates the estimated taxi time for the last aircraft in each runway queue.

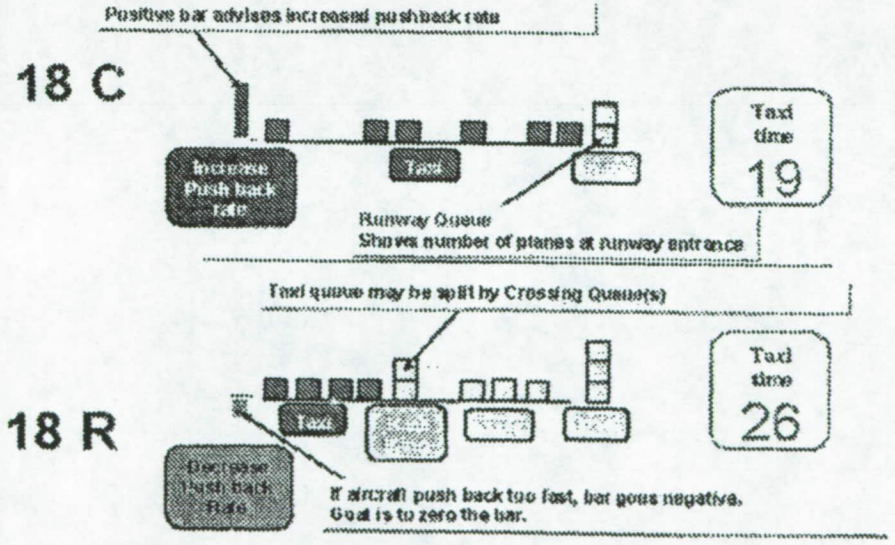

Figure 2. Nominal display for an airport with dual departure runways

Equalizing the taxi times for the last aircraft in the queues balances the runways despite differences in their overall taxi path lengths or differences in their departure capacities. Runway capacity differences can be caused by departure airspace constraints or by differences in mode of operation. For example, a runway dedicated to departures has greater departure capacity than one shared by arrivals. Balancing the taxi times assures that the flow to one runway does not dry up earlier than the flow to the other at the end of a departure rush. The simple taxi-out aid does not attempt to balance taxi delays. Balancing taxi delays (where delay for each runway is reckoned relative to its minimum taxi-out time) could assign too many aircraft to the runway with the longer taxi path.

USING THE TAXI AID

Large hubbing airports with multiple departure runways often allow airline gate/ramp control personnel to control the pushback and ramp taxi process for their own aircraft. They interact with the FAA Clearance Control, Ground Control, and Local Control positions.

After FAA Clearance Control has issued a flight plan clearance and handed off the flight strip to FAA Ground Control, the next step in the departure process occurs when the pilot notifies the carrier's Gate/Ramp Control that he is ready to push back. Gate/Ramp Control issues the pushback clearance and clears the aircraft to the desired taxiway entrance spot. Upon reaching the spot, the pilot monitors the Ground Control frequency in anticipation of a taxi clearance.

Ground Control positively identifies each aircraft that is ready to enter the surface movement area and, when ready, clears it to enter the taxiway. Ground Control handles the aircraft until it first approaches an active runway. Transfer to FAA Local Control may take place before crossing an active runway or before entering the departure runway, whichever occurs first. Ground Control hands off the flight strip to Local Control and authorizes the pilot to switch to the Local Control frequency. Local Control then issues all remaining clearances including the takeoff clearance.

If two or more carriers simultaneously push back aircraft for departure, FAA Ground Control establishes the overall taxi-out order to sequence the aircraft on the runways in a fair and equitable manner subject to departure flow control and wake vortex spacing restrictions. Ground Control also exercises short-term control over the runway balance and the queue lengths for the individual runways by managing taxi out from the ramp area.

Surface surveillance displays will be used in the tower, the Traffic Management Units in the TRACON, the en route Center (where it could help predict near-term departure demand), and the airline operational centers and gate and ramp control positions. The simple taxi-out aid will be used at all the same locations.

In departure pushes FAA tower controllers estimate the size of departure queues by monitoring the occupancy of the flight strip bays and by visually observing the aircraft on the runways and taxiways. The taxi-out aid provides Ground Control with a direct estimate of the current runway balance and suggests the optimum apportionment of aircraft to multiple departure runways. By comparing the queue size to the visual scene or to the tower flight strip bays its users can assure that it accounts for all the aircraft, assigns each aircraft to its proper departure runway, and accurately shows the distribution of the aircraft along the paths to the runways. Users can click on aircraft boxes to display data tags confirming aircraft identity and status.

Gate/Ramp controllers often have difficulty directly observing the details of the other carriers' departure operations or the overall departure demand and capacity. The taxi-out aid provides the needed visibility and suggests near-term limits to the pushback rate for all of the carriers. Individual carriers can then infer their own pushback rate limits by comparing the pushback advisories with their recent departure flow performance.

The simple taxi-out aid facilitates collaborative decision making among the carriers and ATC by providing all participants with common situational awareness. They collaboratively determine the throughput performance objectives for the airport and enter the desired departure rate into the simple taxi-out aid. They can change the control performance goals at any time, resulting in immediate changes to the displayed pushback rate advisories. When a single carrier dominates the departure push the pushback advisories more directly apply to the dominant carrier, simplifying the cooperative management of gate or ramp holds between the hubbing carrier and FAA Ground Control.

THE PUSHBACK ADVISORY ALGORITHM

The taxi-out aid automatically determines the relationship between departure throughput, taxi-out time, and effective queue length for each taxiway/runway path. The effective queue for a runway includes all aircraft on their way to the runway as well as the aircraft in the physical queue that the runway entrance. The relationship between these performance measures is obtained by continually measuring,

tracking, and archiving each of these quantities. The taxi aid analyzes the archived data under all operating conditions to determine and model the needed relationships.

Figure 3 illustrates the relationships used to develop the departure performance database. It relates the overall mean departure throughput to the effective queue length for BOS in August 1991, which includes all aircraft that have pushed back but not yet taken off.

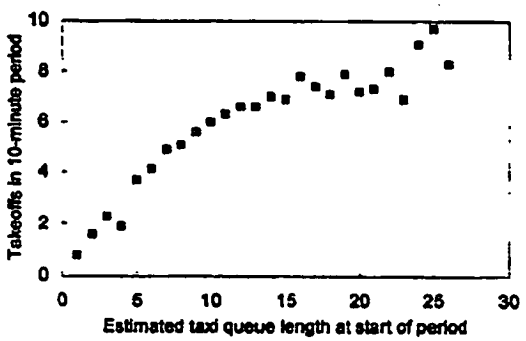

Figure 3. Mean departure rate vs. taxi occupancy for BOS, Aug 1991 - from (Shumsky 1997).[15]

Taxi-out delays exhibit classical queuing behavior. In departure rushes, aircraft fill the taxiways, queues grow at runways, the ambient runway/airspace departure capacity limits the takeoff rate, and delay increases faster than the queue length. The service rate of a queue is always the minimum of the demand and the capacity. Figure 3 shows that when there are less than 15 departure aircraft on the taxiways, there are gaps in the flow to the runway, and the takeoff rate is determined mainly by the departure demand. When there are many more aircraft on the taxiways than 15, there is usually a queue of aircraft at the runway entrances. The takeoff rate is then determined by the departure capacity of the airport. The variability in mean takeoff rate for large queue lengths reflects the variation in airport capacity that occurred during the 1-month data-gathering period. If the operator desires a total throughput of seven aircraft in 10 minutes, this curve tells him that the queue length should be about fifteen. Increasing the queue length above 15 will have little operational effect other than to increase taxi-out delays.

The taxi aid bases the pushback limit calculation on the principle that the rate of growth of the departure queue is determined by the difference between the pushback rate and the runway take-off rate. The taxi aid determines the pushback limit from the queuing conservation relationship:

$$P = N_T - N_0 + D,$$

where P is the number of pushbacks during the next T minutes; N_T is the desired queue length T minutes in the future; N_0 is the present queue length, and D is the number of takeoffs expected during the next T minutes.

The desired future queue length, N_T is obtained from the performance database after a throughput goal has been established by the users. The present queue length, N_0 is obtained from surface surveillance. The number of takeoffs during the next T minutes is obtained by predicting the departure rate over the next T minutes.

To predict the departure rate the simple departure aid uses historical data relating queue length to departure rate. The departure rate is determined by the departure capacity of the airport when the departure queue is large. In a departure push the departure capacity varies slowly and can be reliably tracked and predicted in real time with no knowledge of the weather, the runway configuration, or the arrival/departure mix.[15] Shumsky was able to forecast the takeoff rate 30 minutes into the future with a RMS error of about two departures per hour. He achieved his best capacity estimates by fitting an analytical approximation to data of the type shown in Fig. 3. He used an exponential curve fit to provide a static model of the relationship between queue length and departure rate. He then added a dynamic term to adjust the curve up or down at 10-minute intervals. The dynamic adjustment linearly tracked the observed takeoff rate by smoothing over the residuals of the takeoff estimate at each update interval.

Shumsky did not have surface surveillance data to help refine his estimate. Knowing the runway configuration, the distribution of aircraft on the surface, and the balance between arrivals and departures for each runway makes it possible to improve the static estimates by modeling families of curves for different operating conditions. Tracking repeatable daily or hourly variations can further improve the estimation process.

The taxi aid also estimates the total taxi-out time for the next pushback to help balance the queues. The relationship between effective queue length and taxi time is also obtained from the performance database. This is illustrated in Figure 4, which relates the mean taxi-out time to the queue length for the entire departure taxiway system at Boston (BOS) from January to March 1997. This curve indicates that a queue length of 15 aircraft produces an expected taxi-out time of 36 minutes. The mean taxi-out delay relative to the unimpeded delay can also be obtained from this figure:

the taxi delay is approximately 36 - 12 = 24 minutes for a queue length of 15 aircraft.

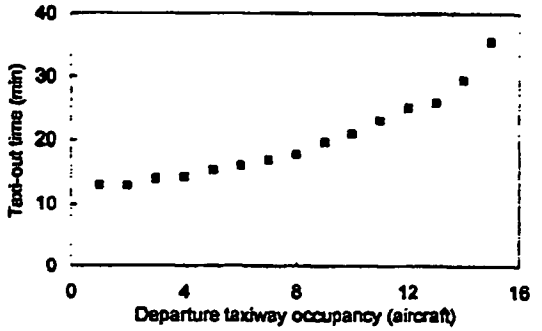

Figure 4. Dependence of mean taxi out time on the aircraft count in the BOS departure taxiway system at the time of pushback - from (Idris et al, 1999).[12]

To estimate the departure queue length in real time for a multi-runway airport, the taxi aid associates each aircraft with a departure runway and automatically detects changes in departure runway assignment. At pushback the taxi aid estimates the departure runway for each aircraft from the filed flight plan for the aircraft and the current runway configuration. The taxi aid checks the position of the aircraft at each subsequent update against a site-adapted list of departure taxi routes for that configuration to confirm that the aircraft is still taxiing towards the same runway.

OTHER DATA SOURCES

Although the simple taxi-out aid needs only flight plan data and surface surveillance data to function, additional data types would be valuable for surface movement predictions. Pre-departure pushback status information and automatic pushback notices from all aircraft operators at the airport, final approach surveillance data from ARTS, CTAS data, and data on operational constraints in departure airspace would all improve the performance of taxi aids.

These data sources would help extend departure demand predictions farther into the future to better support strategic planning. Pre-departure status information from aircraft operators would be very useful if it were reliably available for all departures. Taxi aids will require timely automatic pushback data at an airport whose surface surveillance does not provide reliable coverage in the ramp area. It is essential to measuring the gate and ramp performance of the surface surveillance system as the first step in implementing a taxi aid at an airport.

CTAS data includes arrival planning information as well as ARTS approach surveillance data. These would both be necessary for more advanced taxi aids that manage arrivals as well as departures.

Other information would be useful in the future to refine the departure capacity estimates. Aircraft type information (which is included in the flight plan) would help determine wake vortex spacing requirements on take-off.[7] Although it is possible to automatically detect runway configuration changes in near real time from either ARTS surveillance data or surface surveillance data knowledge of planned configuration changes would obviously improve the predictability of departure capacity.[8] Finally, information on current and planned departure airspace restrictions would help improve both the capacity estimates and the archived data used for modeling constraints on departure sequencing and timing.

BENEFIT MECHANISMS

The principal benefit expected from the simple taxi aid is a reduction in direct operating costs for aircraft operators from reduced taxi-out delay. The aid will also reduce environmental pollution and controller workload. It will achieve these benefits by helping controllers and aircraft operators determine when to use gate and ramp holds to avoid overloading departure queues at individual runways. At airports with two or more departure runways, it will help reduce departure delays by balancing the queues to prevent under-utilization of one runway while overloading another. Predicting such unbalances early will help reduce expensive delays for aircraft that are not yet committed to depart on a particular runway.[13] However, when runways become unbalanced because of procedures that rigidly map departure fixes to runways, the solution will likely involve procedural changes to allow an aircraft to depart from a runway not normally associated with its planned departure fix.

In addition to avoiding runway saturation and runway imbalances, it also helps ground controllers plan taxi clearances and assign departure runways in order to provide continuous streams of traffic for all departure runways. The magnitude of these benefits can be determined initially by analysis and then by using controller-in-the-loop simulations and re-enacting demanding departure taxi scenarios. Baseline tests can be run on the scenarios to quantify the resulting departure delays with and without the help of pushback advisories.

Although the initial functionality will be limited to pushback advisories and departure queue length

predictions, its use of surface surveillance and digital flight plan data to predict surface movements will be a necessary first step in enabling more efficient management of surface movements and in supporting future functionality. Future efficiency-enhancing functionality based on the simple departure aid could:

- help coordinate arrival and departure planning by accurately predicting departure times,
- estimate taxi-in as well as taxi-out delays for individual aircraft,
- help provide runway assignment and sequencing advisories for arrivals as well as departures,
- help support procedures that permit an aircraft to depart from a runway not normally associated with the departure fix designated in its flight plan,
- help reduce runway crossing delay by reducing taxi time variances and excess buffers,
- help improve reconfiguration planning,
- help increase landing throughput by monitoring runway occupancy times, and
- help reduce de-icing delays by monitoring de-icing queues and time violations.

All of these efficiency enhancements will require the runway adaptation, surveillance and flight plan processing, surveillance data analysis, capacity prediction, and display generation capabilities that are required in a the simple taxi-out aid. These capabilities will also provide the basis for important surface safety functionality.

RISKS

The fundamental risk inherent in any activity to develop decision support tools is the failure to identify the main operational problems and the consequential failure to deliver dollar, workload, or environmental benefits at the selected site. Data is needed to verify assumptions about delay and cost mechanisms at candidate sites. The surface surveillance system and FAA delay reporting systems must be used along with operational logs and consultation with controllers and aircraft operators to determine the principal causes of taxi-out delay.

Failure to achieve user acceptance is a fundamental risk. The controller interface design and development process involves significant research risk. Familiarization, training, and display modifications will be required, and simulations will be needed with and without the pushback advisories and queue distribution graphics to determine if they are useful and reduce workload for the users.

It will also be necessary to find space for a new display in the tower. Experiments with trial interface designs must begin as soon as possible in a tower environment. The NASA tower cab simulator is ideal for this purpose. Adapting the NASA simulator to the chosen airport should begin as soon as the test site has been chosen. This adaptation is expensive, so there are budgetary risks involved in this site decision. However, evaluation of generic displays could begin immediately with a tower cab simulation of any airport.

Major research risks were noted above in discussing functionality and data interfaces. An important risk is the performance and accuracy of the key capacity tracking and prediction algorithms used to estimate the future departure queue length. These estimates will rely on processing and tracking routines that continuously monitor, record, and analyze operational surface surveillance at each airport. Although data tracking is performed routinely in Air Traffic control surveillance systems, it is not commonly done to automatically obtain and update operational information needed for decision support automation.[3,10] If it is not possible to predict departure capacity with sufficient accuracy solely from recent observations of departure queues and takeoff rates, it may be necessary to obtain additional sources of data.

Important programmatic and technical risks involve access to, continued availability of, and coverage of the surface surveillance data. Surveillance data is key to automatically determining the current queuing status for each departure runway and automatically predicting near-term departure demand. The FAA is developing a commercial surface multilateration system intended to provide reliable surveillance and identification of all aircraft with operating transponders on airport movement areas. The data is fused with data from surface surveillance radars, with data from the ARTS terminal airborne surveillance system, and with flight plan data.

The FAA has initiated development of an operationally deployable digital surface surveillance system known as ASDE-X. This program is intended to lead to an operational capability that includes multilateration in conjunction with a low-cost primary radar operating at X-band.

A minor regulatory change is needed to obtain reliable surveillance and identification coverage for all aircraft in the ramp and movement areas. Transponder multilateration, which is the key to low-cost, clutter-free surveillance on the surface, depends on transmissions from aircraft transponders. Official FAA procedures must change at airports with multilateration

systems to mandate that aircraft on the airport ramp and movement areas leave their transponders operating at all times rather than turn them off as currently required. Because there appears to be no technical or interference problem from continuous operation of transponders on the surface of major airports there is no technical or operational impediment preventing this regulatory change.[20]

There is a small technical risk that surface surveillance may not provide reliable coverage in ramp and gate areas because of blockage of transponder transmissions by airport structures or by the airframes of other aircraft. Recent tests at DFW addressed the coverage issue.[21] The gate coverage appeared good based on limited operation within the ramp and gate areas. Multilateration coverage can always be improved by the addition of additional ground sensors. In addition, a few aircraft will not be visible because they have been intentionally wired to reduce controller workload by automatically switching off their transponders when they are on the ground. One of the early research activities in any program to develop surveillance-based taxi aids must be a complete characterization of the coverage issue.

REFERENCES

1. Anagnostakis, Idris, Clarke, Feron, Hansman, Odoni, and Hall (2000) "A Conceptual Design of a Departure Planner Decision Aid," 3rd USA/Europe Air Traffic Management R&D Seminar, Napoli.

2. Anderson, Carr, Feron, and Hall (2000), "Analysis and Modeling of Ground Operations at Hub Airports," 3rd USA/Europe Air Traffic Management R&D Seminar, Napoli.

3. Andrews (2001), "Radar-Based Analysis of the Efficiency of Runway Use," AIAA GNC Conf., Montreal.

4. ATA (2000), "Approaching Gridlock" Air Transport Assocation www.airlines.org/public/gridlock/ag2a.htm

5. Bhaumik et al (1996), "Surface Movement Advisor SMA Operations Manual," SMA Project Office NASA Ames Research Center Moffett Field, California, SMA-152, June 1, 1996.

6. CODAS (2000) "Federal Aviation Administration Office of Aviation Policy and Plans, Consolidated Operations and Delay Analysis System," www.apo.data.faa.gov.

7. Dasey (1998), "A Departure Wake Vortex Monitoring System: Concept and Benefits", 2nd USA/Europe Air Traffic Management R&D Seminar, Orlando.

8. Eggert (1994), "Demonstration of Runway Status Lights at Logan Airport", Linc. Lab. J. 7, 169.

9. Edelman (2000), "Surface Management," Briefing to Surface R&D Technical Interchange Meeting, 7/11/2000, MIT, Cambridge, MA

10. Erickson and Prusak (1994), "Flight Time Modifier Flight Time Data Modifier," Proceedings of the National Traffic Management Conference, Reston VA 1/24/94.

11. Erzberger (1990), "Integrated Automation Tools for the Center and the TRACON", The Journal of ATC, 32(4).

12. Idris, Anagnostakis, Delcaire, Hansman, Clarke, Feron, and Odoni (1999), "Observations of Departure Processes at Logan Airport to Support the Development of Departure Planning Tools," Air Traffic Control Quarterly, 7(4).

13. Lawson (1998), "Surface Movement Advisor (SMA)", J. Air Traffic Control, January-March, 1998.

14. Pujet, Delcaire, and Feron (2000) "Input-Output Modeling and Control of the Departure Process of Congested Airports," Air Traffic Control Q., 8(1).

15. Shumsky (1997), "Real Time Forecasts of Aircraft Departure Queues," Air Traffic Control Quarterly, 5(4).

16. Signor (1999), "Operations Concept for the Surface Management System (SMS) Prototype," NASA Ames Research Center Aviation Surface Technologies Area Team Report SMS-102, Moffett Field, CA.

17. Signor (2000), "Surface Management System Research & Development Plan Volume I – Technical Plan," NASA Ames Research Center Aviation Surface Technologies Area Team Report SMS-101, Moffett Field, CA.

18. Vail (2000), "Operational Concept for an Airport Surface Management System (ASMS)", Briefing to Surface R&D Technical Interchange Meeting, 7/11/2000, MIT, Cambridge, MA.

19. Wilhelmsen (1994), "Preventing Runway Conflicts: The Role of Airport Surveillance, Tower-Cab alerts, and Runway Status Lights," Linc. Lab. J. 7, 149.

20. Wood and Bush (1998), "Multilateration on Mode S and ATCRBS Signals at Atlanta's Hartsfield Airport." MIT Lincoln Laboratory Project Report ATC-260, available to the public through NTIS, Springfield, VA 22161.

21. Wood (2000), "Multilateration System Development History and Performance at Dallas/Ft Worth Airport," Proceedings, 19th Digital Avionics Systems Conference, October 2000, Philadelphia.

LIST OF ACRONYMS

ACM	Airport Configuration Manager
ARTCC	Air Route Traffic Control Center
ASR-9	Airport Surveillance Radar
ATC	Air Traffic Control
CONUS	Continental United States
CTAS	Center/TRACON Automation System
DFW	Dallas/Ft. Worth International Airport
ETA	Estimated Time of Arrival
FAA	Federal Aviation Administration
FAST	Final Approach Spacing Tool
FPODEV	Forecasted Probability of Deviation
GUI	Graphical User Interface
ITWS	Integrated Terminal Weather System
MIT/LL	Massachusetts Institute of Technology Lincoln Laboratory
NEXRAD	Next Generation Weather Radar
OTWD	Offline Traffic and Weather Display
P3I	Pre-Planned Product Improvement
PGUI	Planview Display Graphical User Interface
PODEV	Probability of Deviation
RUC	Rapid Update Cycle
TCWF	Terminal Convective Weather Forecast
TDWR	Terminal Doppler Weather Radar
TMC	Traffic Management Coordinator
TRACON	Terminal Radar Approach Control
VIP	Video Integrator and Processor

TECHNICAL REPORT STANDARD TITLE PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NASA/A-5		

4. Title and Subtitle	5. Report Date
Contributions to the AIAA Guidance, Navigation & Control Conference	23 January 2002
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Steven D. Campbell, Editor	NASA/A-5

9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02420-9108	11. Contract or Grant No. NASA Ames

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
NASA Ames Research Center Moffett Field, CA 94035 Federal Aviation Administration Washington, DC 20591	Project Report
	14. Sponsoring Agency Code

15. Supplementary Notes

This report is based on studies performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology, under Air Force Contract F19628-00-C-0002.

16. Abstract

This report contains six papers presented by the Lincoln Laboratory Air Traffic Control Systems Group at the American Institute of Aeronautics & Astronautics (AIAA) Guidance, Navigation and Control (GNC) conference on 6-9 August 2001 in Montreal, Canada. The work reported was sponsored by the NASA Advanced Air Transportation Technologies (AATT) program and the FAA Free Flight Phase 1 (FFP1) program. The papers are based on studies completed at Lincoln Laboratory in collaboration with staff at NASA Ames Research Center.

These papers were presented in the Air Traffic Automation Session of the conference and fall into three major areas: Traffic Analysis & Benefits Studies, Weather/Automation Integration, and Surface Surveillance. In the first area, a paper by Andrews & Robinson presents an analysis of the efficiency of runway operations at Dallas/Ft. Worth using a tool called PARO, and a paper by Welch, Andrews, & Robinson presents delay benefit results for the Final Approach Spacing Tool (FAST). In the second area, a paper by Campbell, et al. describes a new weather distribution system for the Center/TRACON Automation System (CTAS) that allows ingestion of multiple weather sources, and a paper by Vandevenne, Lloyd, & Hogaboom describes the use of the NOAA Eta model as a backup wind data source for CTAS. Also in this area, a paper by Murphy & Campbell presents initial steps towards integrating weather-impacted routes into FAST. In the third area, a paper by Welch, Bussolari, and Atkins presents an initial operational concept for using surface surveillance to reduce taxi delays.

17. Key Words	18. Distribution Statement
	This document is available to the public through the National Technical Information Service, Springfield, VA 22161.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	80	

FORM DOT F 1700.7 (8-72) Reproduction of completed page authorized

***** General Data *****
039h (1): 20020524
039j (1): STI
040a (1): NASA
090f (1): 2002078359
***** Evaluation Data *****
017c (1): No Copyright
037j (1): Y
0722 (1): NASA Scope and Coverage
072a (1): 04
072b (1): Aircraft Communications and Navigation
090d (1): N
090e (1): 1
2459 (1): Unclassified
355a (1): Unclassified
355e (1): NTIS
359a (1): Unrestricted - Publicly Available
359b (1): NTIS
509a (1): Conference Proceedings
541c (1): Regular
541d (1): 20020524
***** Cataloging Data *****
041a (1): English
088a (1): PB2002-104359
245a (1): Contributions to the AIAA Guidance, Navigation & Control Conference
260c (1): 20020123
260d (1): Jan. 23, 2002
300a (1): 88
541f (1): 0215
541g (1): NTIS
541o (1): Microfiche
700a (3): Campbell, S. D.
700e (3): Author
710a (1): Massachusetts Inst. of Tech., Lexington. Lincoln Lab
710a (2): Massachusetts Inst. of Tech., Lexington. Lincoln Lab.
*National Aeronautics and Space Ad
710e (1): Organizational Source
710e (2): Financial Sponsor
937a (1): CASI
937a (2): CASI
937b (1): Hardcopy
937b (2): Microfiche
937d (1): A05
937d (2): A01
hur (1): 1
mur (1): 1
nur (1): 0