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

AIAA Guidance, Navigation, and Control Conference
6-9 August 2001
Montreal, Quebec, Canada
Technical Area: GN&C Concepts for Air Traffic Control Systems

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

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.

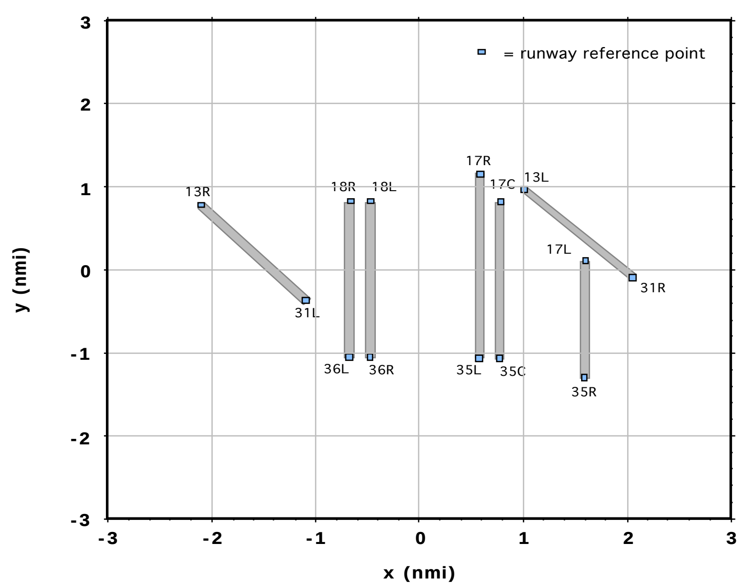

Figure 1. Runway layout at Dallas/Ft. Worth International Airport.

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 nmi. 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.

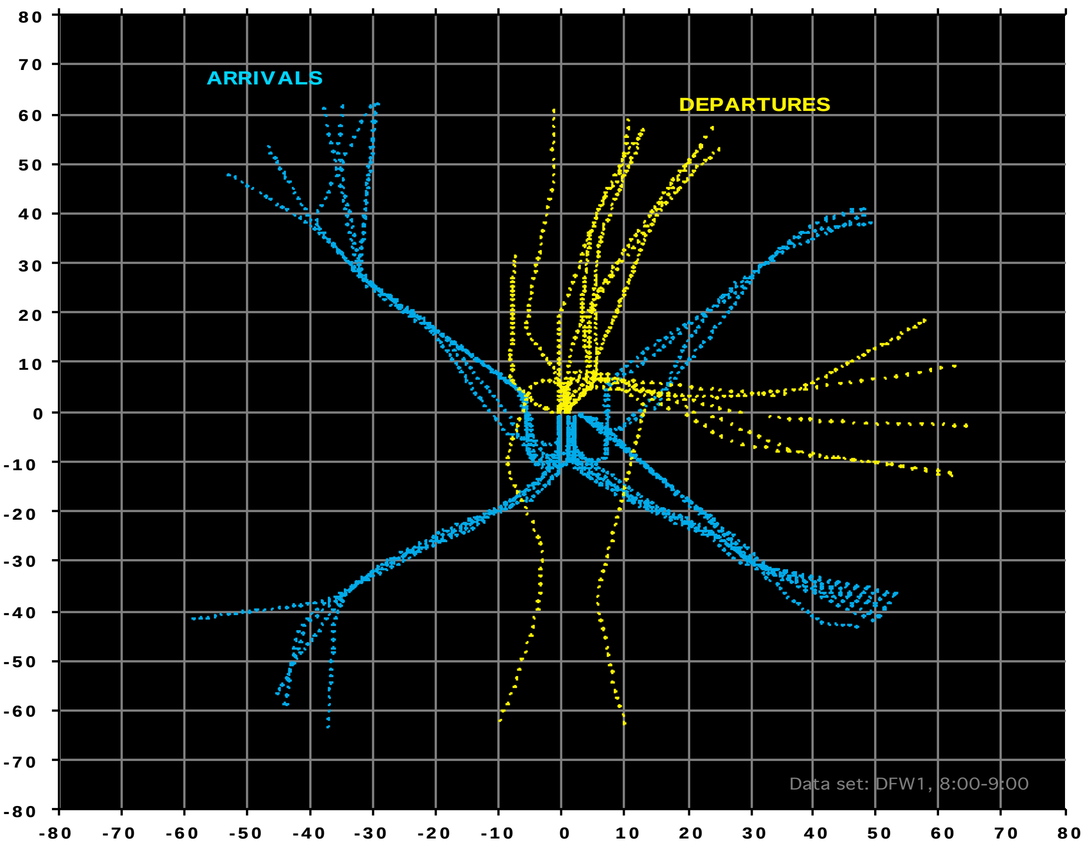

Figure 2. Traffic flow sample at DFW for data set DFW.01.

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	17:59
DFW.04	10 FEB 00	1431	8:19

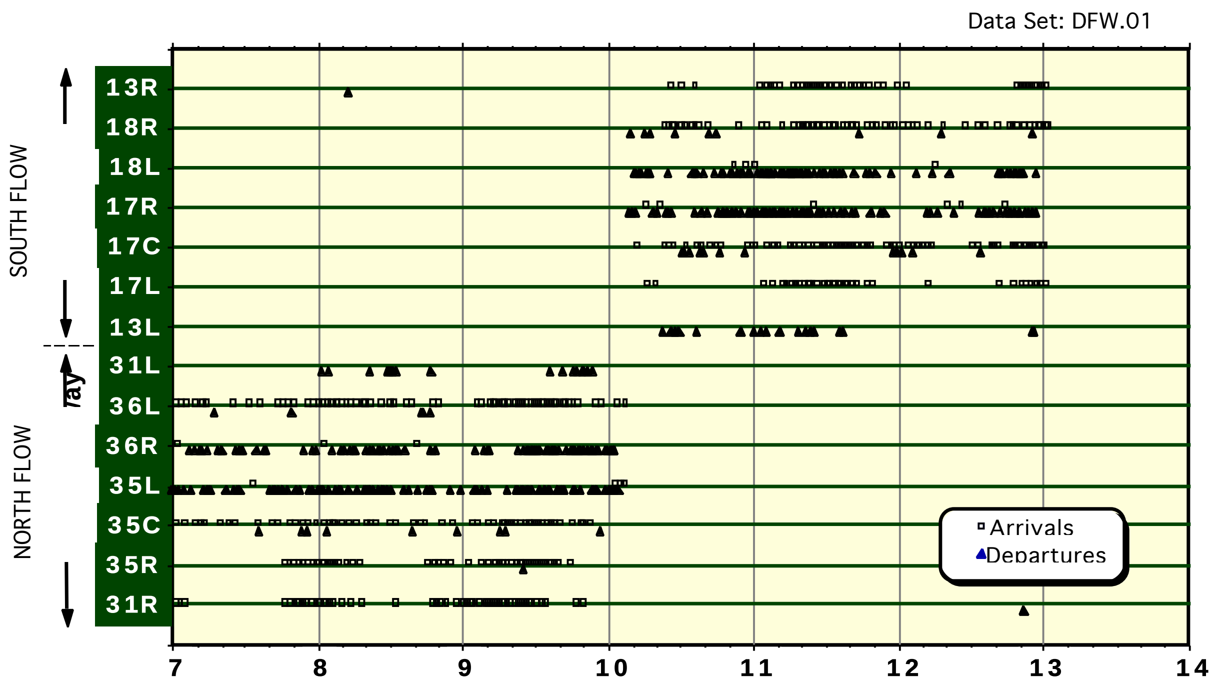

Figure 3. Runway utilization showing a change in runway configuration for data set DFW.01 at 10:30 local time.

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.

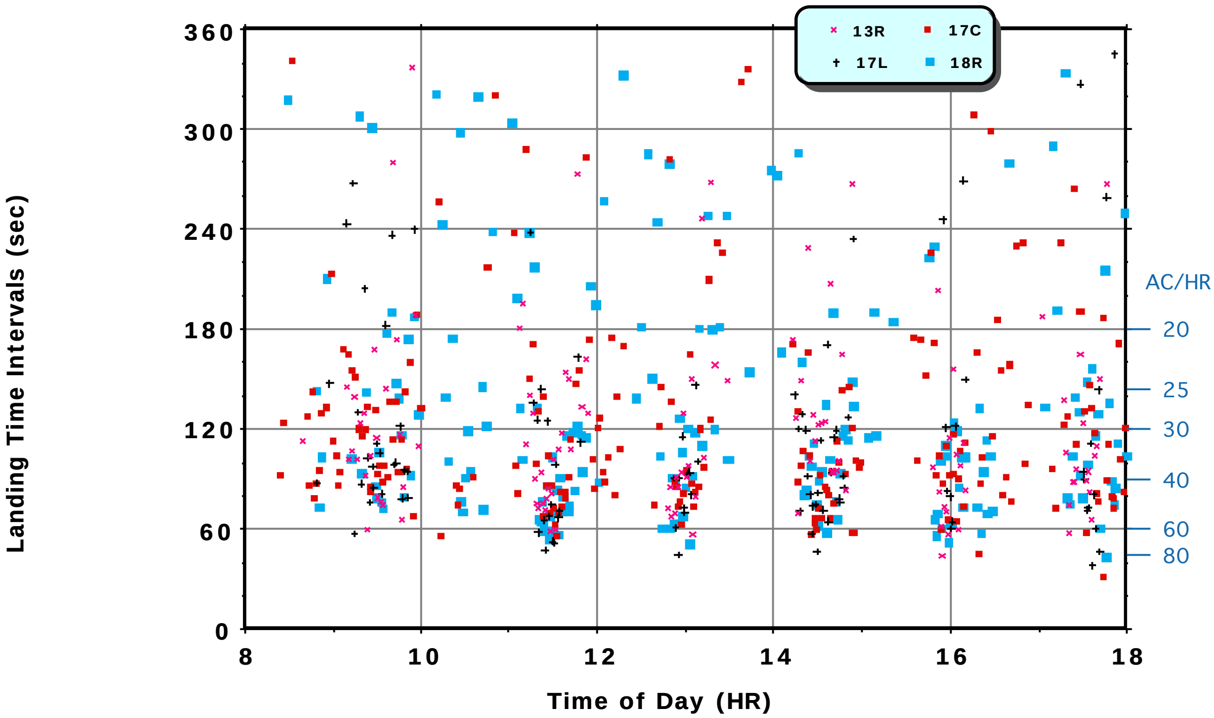

Figure 4. Landing time intervals for data set DFW.04

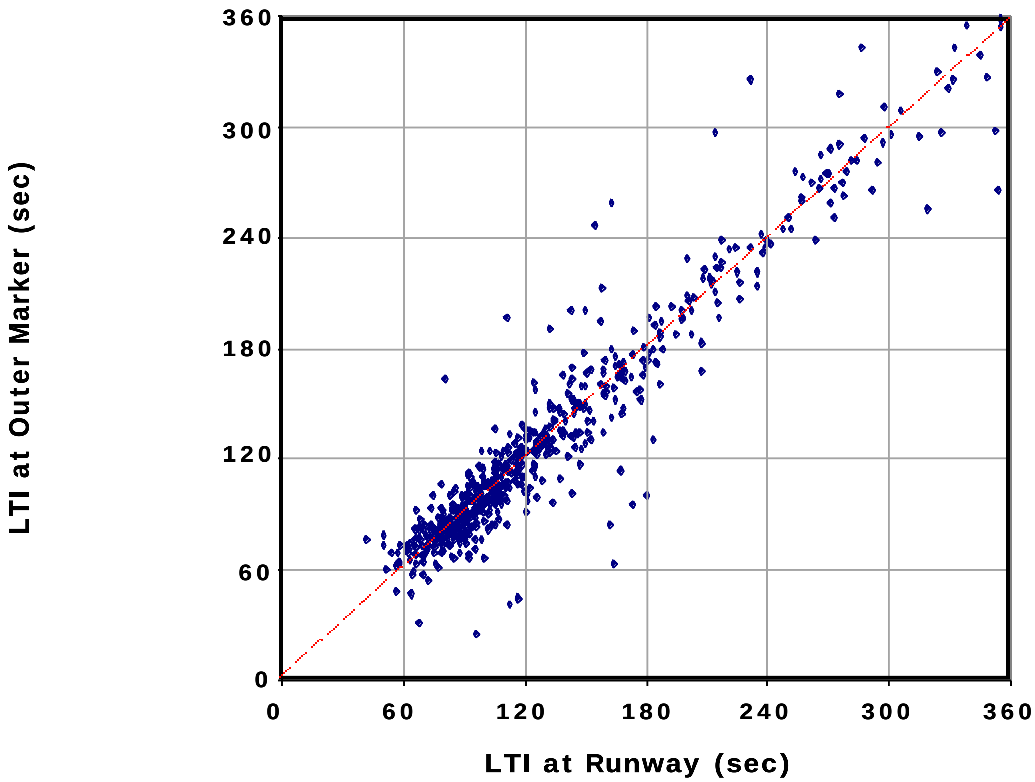

Figure 5. Comparison of LTIs measured at outer marker and at the runway (data set DFW.01).

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.)

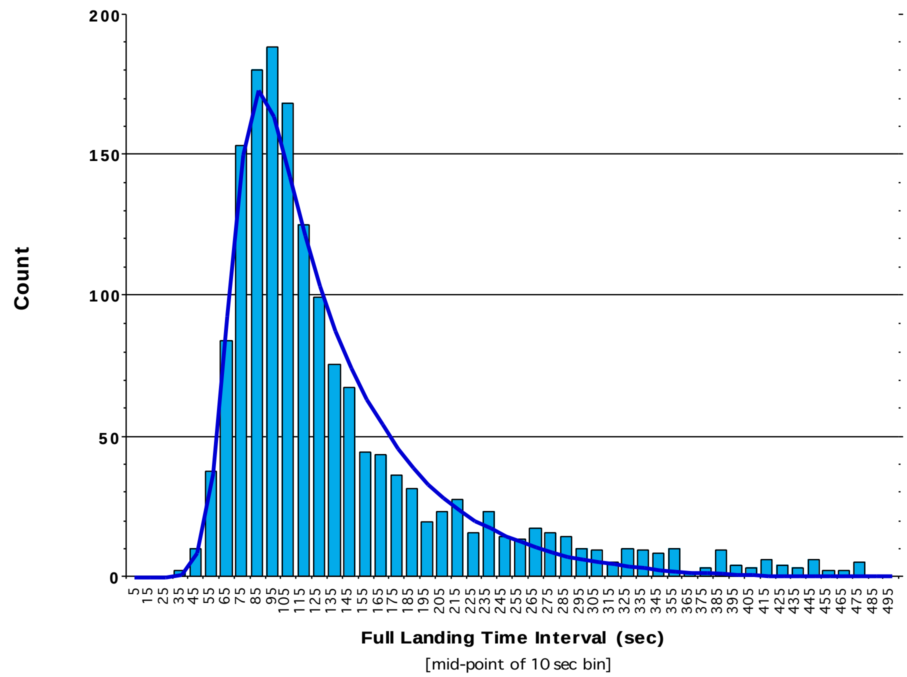

Figure 6. Landing time interval histogram for combined DFW arrivals (1758 pairs).

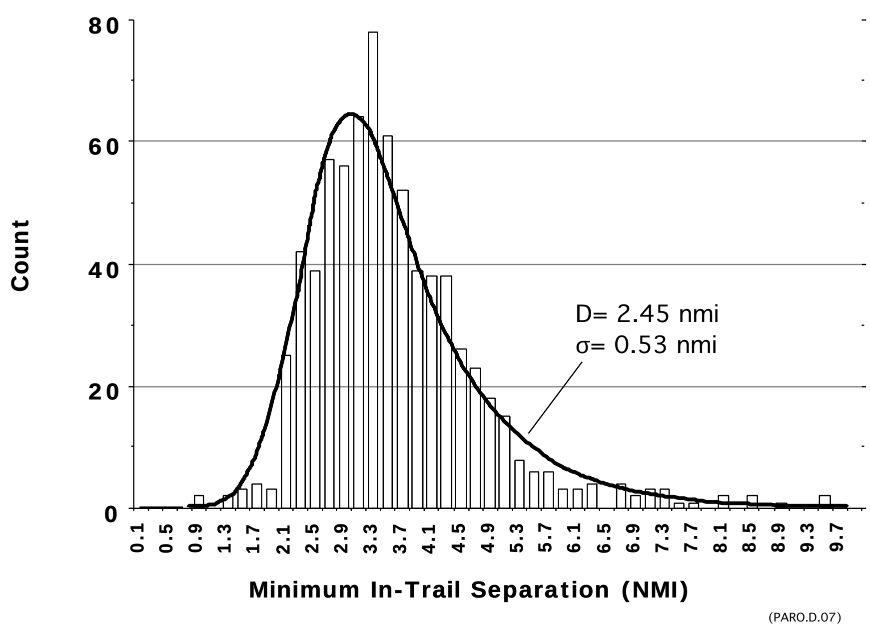

Figure 7. Minimum in-trail separation observed for 801 pairs with both aircraft in the large weight class (DFW).

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

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

where D_0 is the required minimum time separation, F_{SN}^{-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.

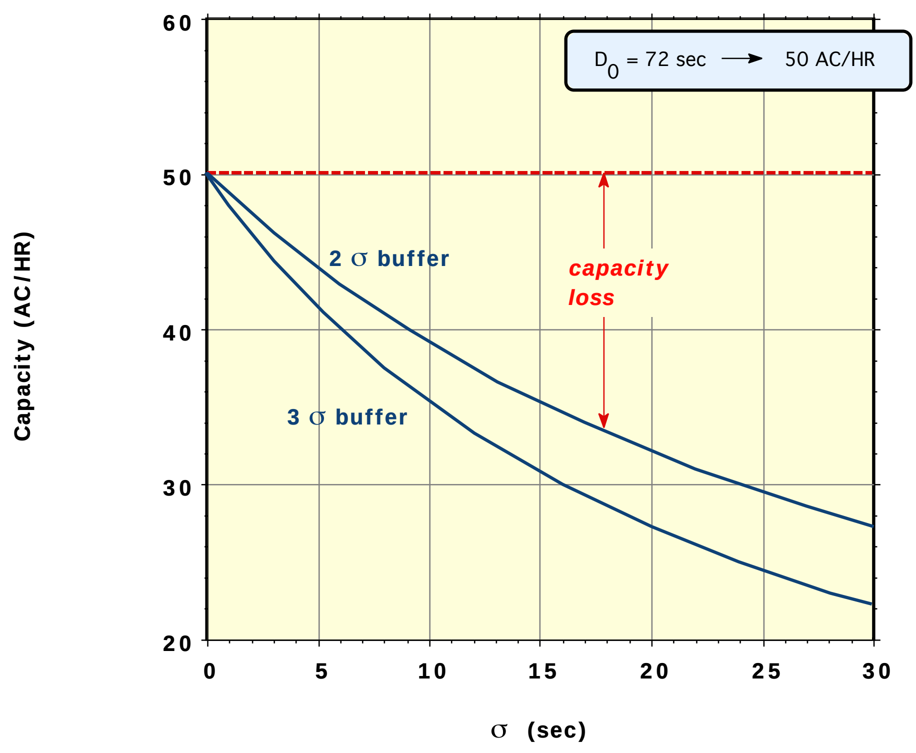

Figure 8. Effect of imprecision (σ) upon runway capacity.

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

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

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_j. Let the mid-point of the separation interval for bin i be $\bar{y}_i$. Then the approximate likelihood function can be written

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

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

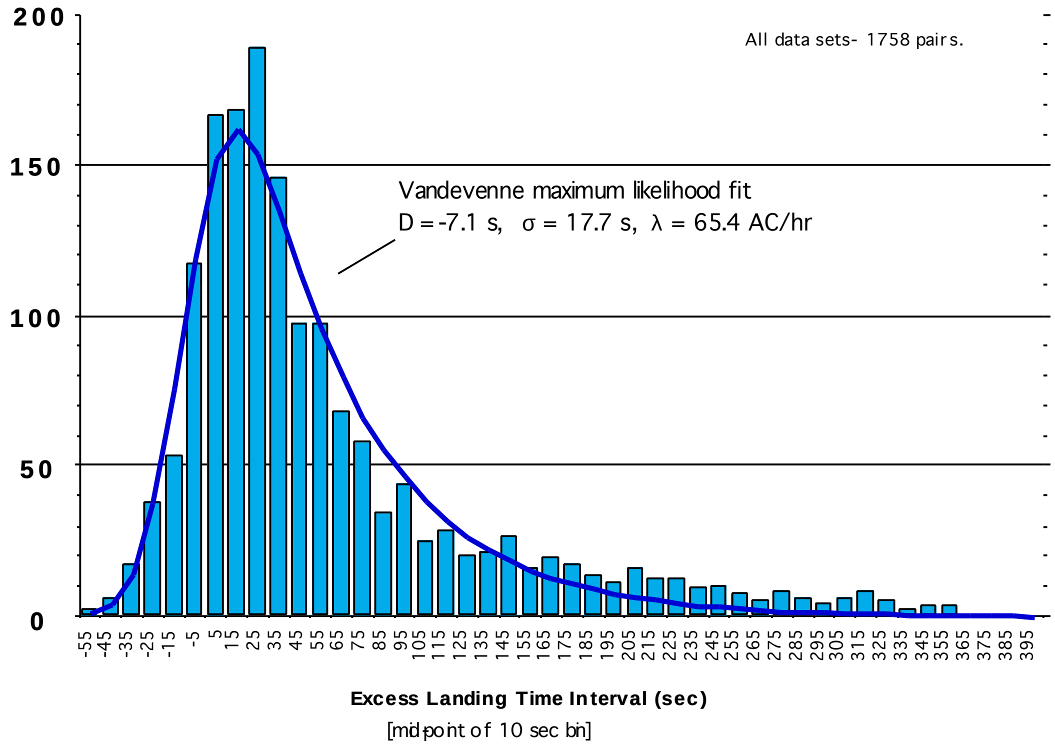

Figure 9. Distribution of excess interarrival spacings (S*) for 1758 DFW arrivals.

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 ρ 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

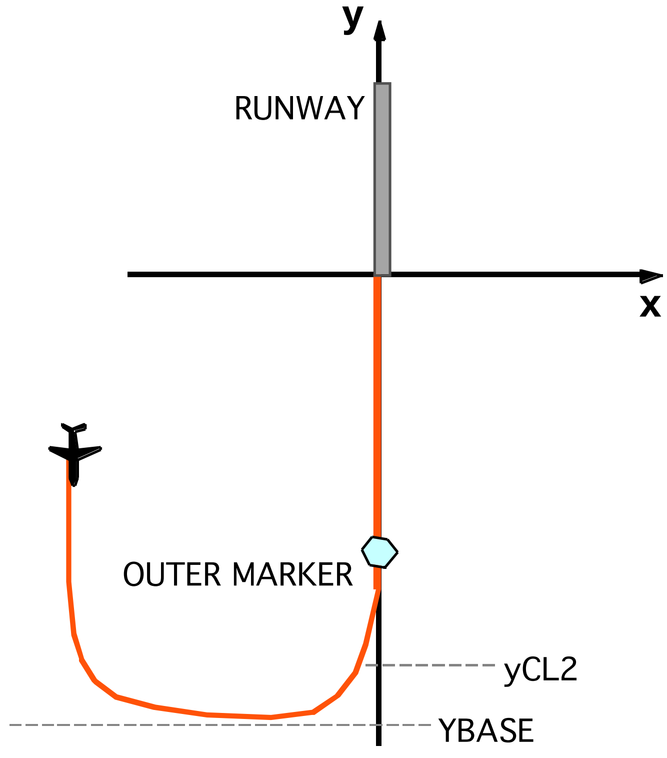

Figure 10. Definition of approach geometry using YBASE and yCL2.

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	•••

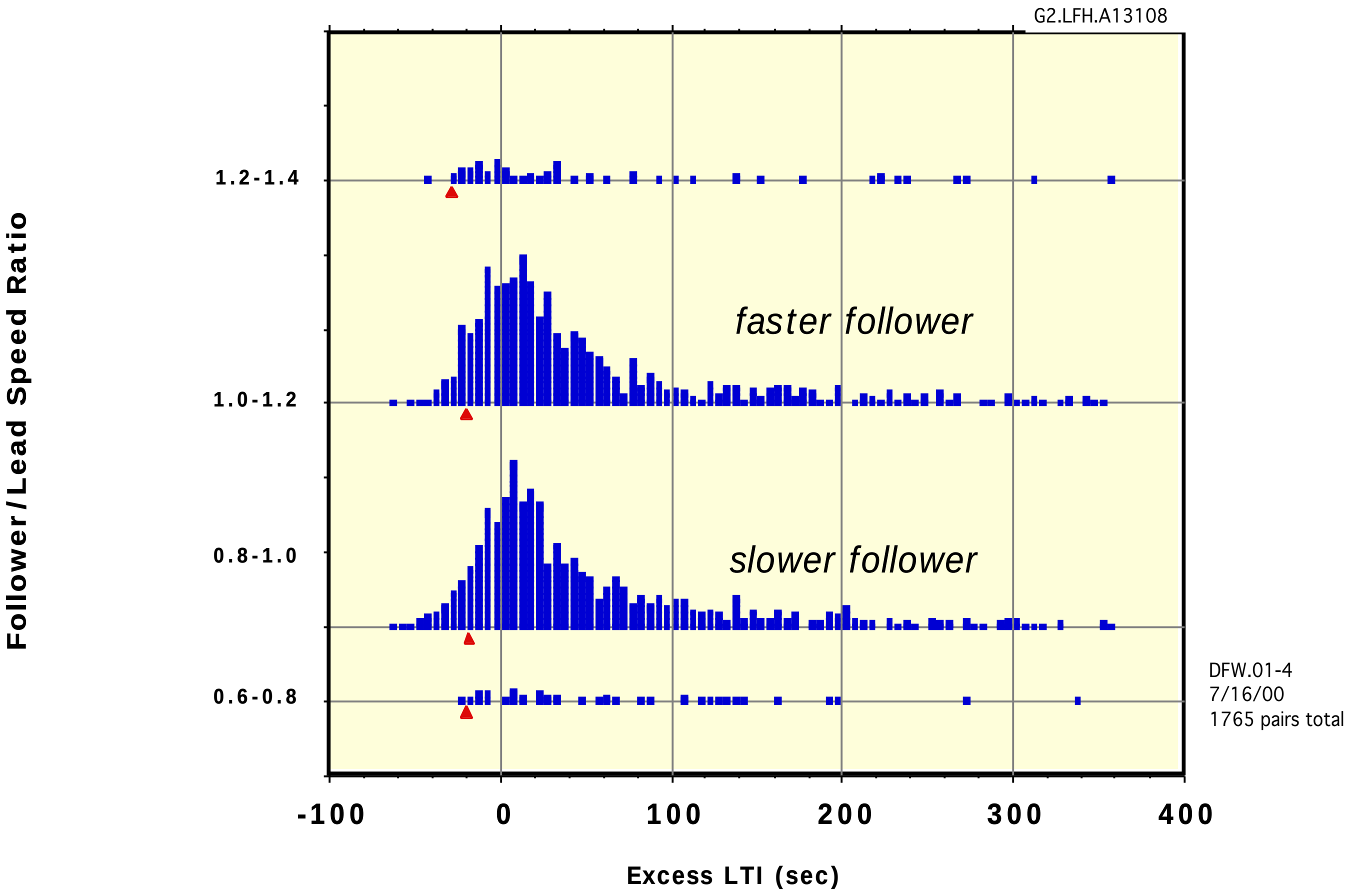

Figure 11. LTI differences between runways

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* distribution for this data, it does appear that aircraft following a downwind leader tend to have higher median values.

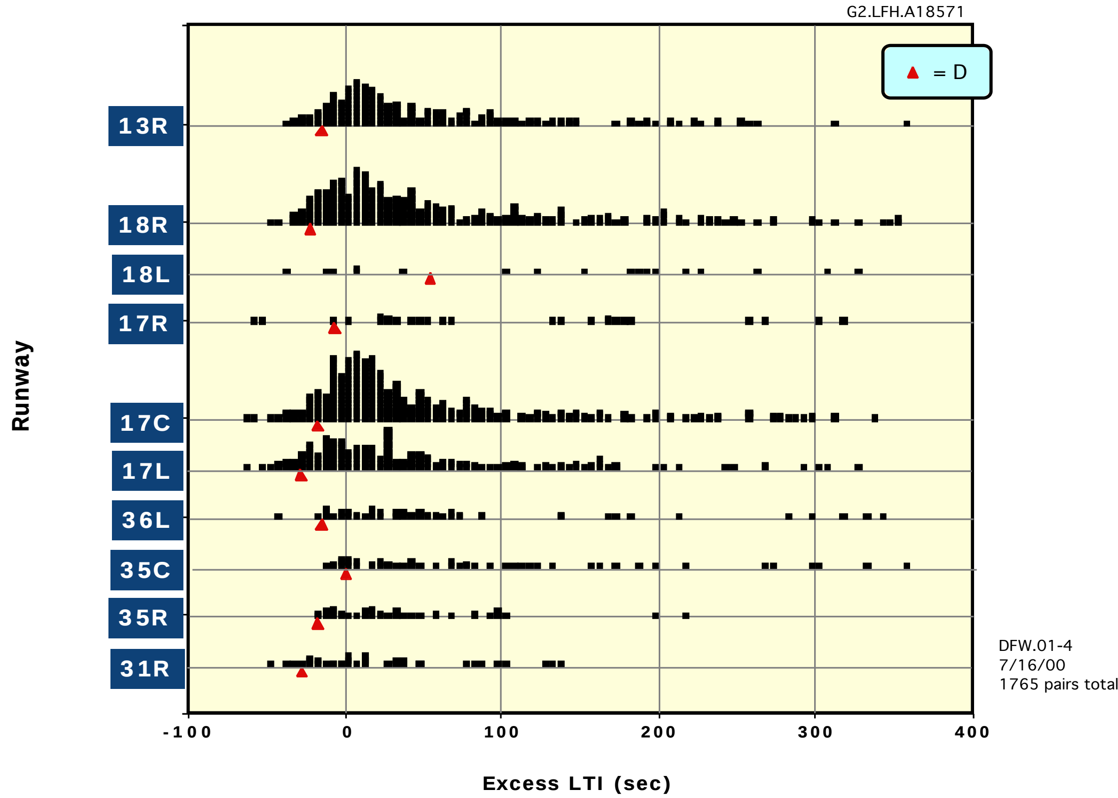

Figure 12. Effect of follower/leader speed ratio on landing time intervals

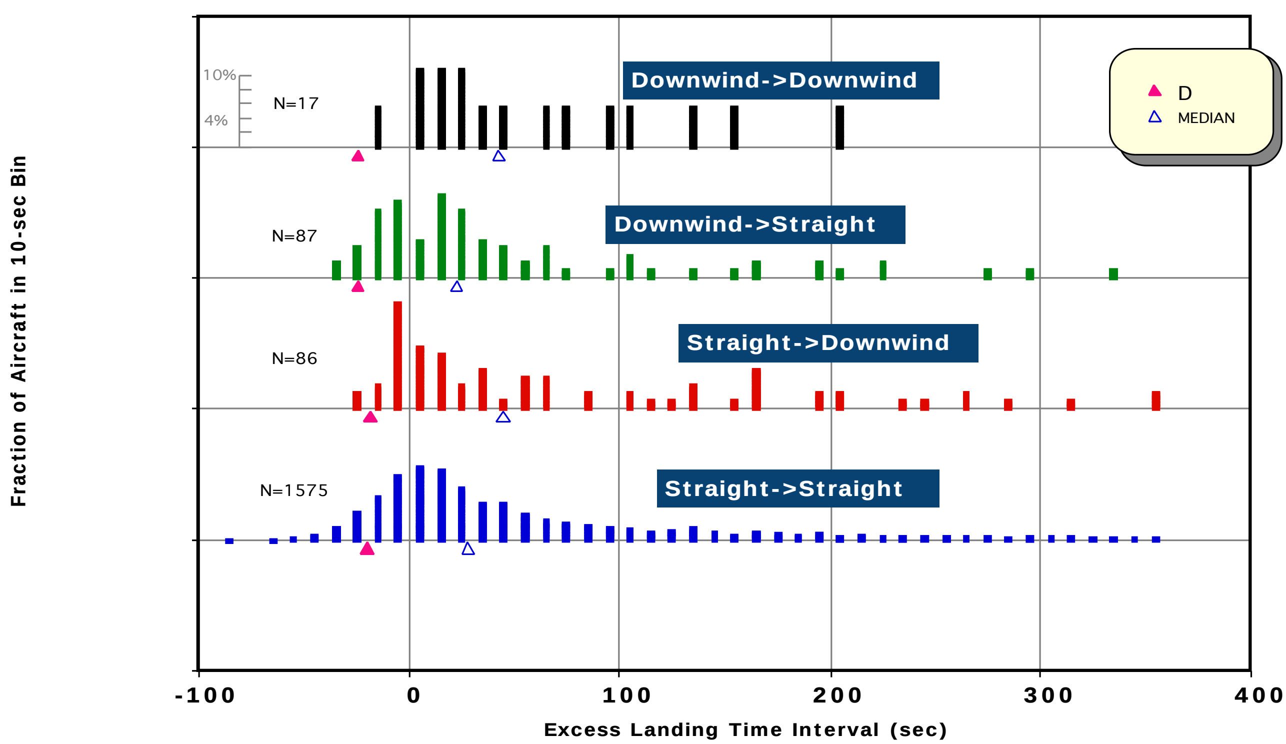

Figure 13. Excess landing time for different approach pair types

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.

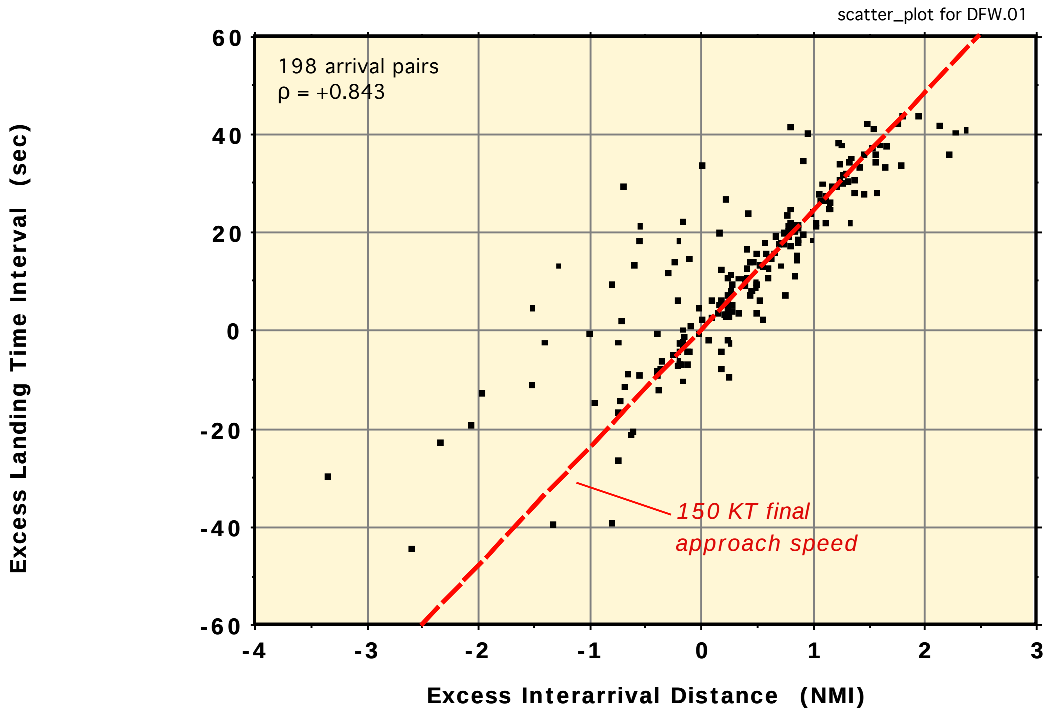

Figure 14. Comparison of excess spacing expressed in time and distance

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].