Normalization of Airport and Terminal Area Operational Performance: A Case Study of Los Angeles International Airport

Mark Hansen and Tatjana Bolic
National Center of Excellence in Aviation Research/Institute of Transportation Studies
University of California, Berkeley
March 2001

1. Introduction

Efforts to improve the performance of aviation infrastructure, whether by deploying mature technologies, introducing new ones, or changing procedures, follow a typical progression. They begin with some initial concept, which is explored, developed, and (when results of the earlier stages are favorable) implemented. Throughout this process, there are repeated efforts to project what the benefits will be. These efforts often involve mathematical models, from simple analytic ones to large-scale simulations. The models predict the effects of implementation on selected performance variables. In using such models, it is natural to make comparisons in which a host of other factors that also affect performance are assumed equal, so that the effects of interest can be isolated.

While implementation is an important milestone in the above process, it should not be viewed as the final step. In particular, ex post evaluation is also necessary in order to determine whether anticipated benefits have actually been realized. The models and assumptions used in pre-deployment assessments are simply not reliable enough to be treated as the ultimate arbiters of truth in this regard. Indeed, ex post analysis, by allowing model predictions to be compared against reality, is an essential means of model validation (or invalidation as the case may be).

But such ex post evaluation faces a significant hurdle. As a given measure is implemented, the world does not stand still. Weather, a major determinant of aviation system performance, is ever changing. Demands on the infrastructure fluctuate over short time scales while growing from year to year. Indeed, such changes in demand often result directly from the change being implemented, as when airlines add flights at peak times in response to new capacity. An additional complication derives from the fact that operations at one airport are affected by conditions at another. If a flight departure is delayed by conditions at the flight origin, this will cause an arrival delay at the destination airport, which may in turn cause further delays at airports downstream. Conversely, a ground delay program may trigger departure delays at an origin airport as a result of congestion at the destination. As a result of these forward and backward propagation effects, performance trends at any one location can be influenced by conditions throughout the aviation system.

Given these circumstances, it is not possible to observe the performance impacts of changes through simple before-and-after comparisons of performance metrics. Rather, it is necessary to normalize such comparisons so that, to the greatest extent possible, other influences can be controlled for. Only when this is done can before-and-after comparisons be translated into with/without ones.

While the need for normalization is easy to recognize, the task of normalization may be quite difficult. There is wide day-to-day and hour-to-hour variation in performance. While many of the sources of this variation are understood, there has been relatively little research that attempts to systematically relate performance variation to its underlying causes. Consequently, we don't know how much of the observed variation in performance results directly from observable differences in weather conditions and fluctuations in demand, and how much is the result of other, less easy to identify, causes. A related question, which has also been little studied, is how to represent and quantify the relationships between performance and the factors influencing it. For example, is it better to treat weather as a set of continuous variables or classify days into weather categories (clear and sunny, stormy, etc)? Or, are the relationships so complex and non-linear as to render normalization impossible by any of these methods?

With such questions in mind, we have undertaken a series of studies in which we statistically analyze performance variation in relation to factors related to weather, demand, and conditions elsewhere in the system. Each of the studies focuses on one airport. The airport studied here is Los Angeles International (LAX). LAX was chosen because it is an implementation site for the FAA's Free Flight Phase I (FFP1) program. Technologies being implemented there include the two Center-Tracon Automation System components: Traffic Management Advisor (TMA), which will be used by the Los Angeles ARTCC, and the Passive Final Approach Spacing Tool (PFAST), which is being implemented at the Southern California TRACON. The immediate purpose of this study, however, is not to assess the benefits from these deployments, but to analyze performance trends at LAX during the pre-deployment period. It is expected that the methodologies developed in this effort may then be used at a later time as part of the FFP1 evaluation.

The essence of our approach is to develop and statistically model a set of daily level performance metrics for LAX. The overarching metric is a weighted average of flight times into LAX, which we term the Daily Flight Time Index (DFTI). The flight time measured is the interval between the scheduled arrival and the actual departure, and can be decomposed into several components, including the departure delay, the taxi-out time, the airborne time, and the taxi-in time. Our research revolves around observing and analyzing day-to-day variation in the DFTI metric and its components.

The remainder of this paper is organized as follows. Section 2 explains the DFTI and summarizes trends in the DFTI at LAX. Sections 3 and 4 present our methods of introducing demand and weather into the normalization process, while Section 5 does the same for conditions elsewhere in the system. In Section 6, we present a baseline linear model of the DFTI. In Section 7, we consider similar models for the individual DFTI components. Section 8 summarizes results for several other model specifications, while Section 9 summarizes our findings.

2. The Daily Flight Time Index

We used individual flight data drawn from the Airline Service Quality Performance (ASQP) database to develop a daily time series of average flight times for flights arriving at LAX. As previously explained, we refer to this daily average value as the Daily Flight Time Index-DFTI. The DFTI is a weighted average of individual flight times into LAX. For individual flights we obtain from ASQP the time interval between *scheduled* departure from the origin gate and *actual* arrival at the gate at LAX. In addition, we obtain the components of this interval, which include the departure delay, the taxi-out time, the airborne time, and the taxi-in time.

As a performance metric, the DFTI is quite similar to average arrival delay. It has two important advantages over that more conventional measure, however. First, it is insensitive to changes in the amount of "padding" built into the flight schedule. As is well known, airlines add extra time in the expectation that their flights will be delayed, and the amount of such padding has generally increased over time. All else equal, the padding increases will cause delays against schedule to go down. By using the DFTI, we avoid this "artificial" effect. Second, use of DFTI permits the decomposition of flight time, and hence delay, changes into the components mentioned above: departure delay, taxi-out time, flight time, and taxi-in time. Since conventional delay metrics require reference to some scheduled time, they cannot be decomposed in this way.

We employ a metric defined at the daily level, rather than some finer time scale, for several reasons. First, as elaborated below, this permits a larger number of flight origins to be included in the averaging. Second, we can treat each day as an independent observation, ignoring interrelationships whereby performance in one time period affects performance in some other period. Such interrelationships cannot be ignored in sub-day time scales, since flights that are unable to land in one time period become additional demand in a subsequent period, often resulting in delays for other flights.

To construct the DFTI (and its components) from the individual flight data, we take a weighted average where the weights reflect the proportion of flights from each origin into LAX over the period of analysis. Since the same weights are used for all the days in the sample, DFTI's are comparable even when the mix of long-haul and short-haul flights changes over time.

We limited the set of origins in the DFTI average in two ways. First, we eliminated airports within 200 miles of LAX. This reduced the influence of correlation between conditions at the origin airport and conditions at LAX. Second, we required that, for every day considered, there must be at least one completed flight from each origin. If an origin has no completed flights to LAX on a given day, one can either exclude the day from the sample or the origin from index. By excluding just a few days, the number of origins that can be included in the DFTI average was greatly increased. In the case of LAX, over the five plus years that we analyzed, we were able to include some 23 origins in the DFTI by eliminating just 4 of 1978 days. The set of origins, and their associated weights, are shown in Table 1. By way of comparison, in order to include 34 origins in the DFTI, it would have been necessary to eliminate 1024 days.

Figures 1, 2, and 3 summarize trends in DFTI for LAX since 1995. Figure 1 simply plots the daily values. It is evident from Figure 1 that DFTI varies substantially from day to day. Most of the time, it is within the range of 140 to 160 minutes. But there are a considerable number of days in which DFTI exceeds 180 minutes, and a few where it goes above 200. Taking 140 minutes as a nominal value for good days with essentially no delayed flights, we see that on very bad days the average delay (measured against the 140 minute standard) can exceed an hour per arrival.

Table 1. Set of Origins and their Associated Weights used for LAX Daily Flight Time Index

Origin Airport	Weight (%)
LAS	12.1
SFO	11.4
PHX	11.1
OAK	9.0
ORD	8.4
SEA	5.9
DFW	5.6
SJC	4.4
PDX	3.8
SMF	3.7
SLC	3.6
HNL	3.5
IAH	3.0
ATL	2.2
TUS	2.2
ABQ	1.8
STL	1.8
MIA	1.7
ELP	1.4
MCO	1.3
PIT	1.3
CVG	1.0

Figure 2 plots the 30-day moving average of the DFTI, with the data points coded by season. A seasonal pattern is evident, with higher values in the fall and winter and the lowest values generally in the summer. Recent years have seen some divergence from this pattern. The springs of the last three years have been considerably worse than previous ones, as was the summer of 1999 compared to those of 1996-1998. Also, during the fall of 1999, the DFTI was stable throughout the quarter instead of climbing toward the end. The early winter of 2000 was also considerably better than average, although high DFTI values returned in the latter part of that quarter. Overall, Figure 2 suggests some trend toward increasing DFTI over the past five years, although that trend is dominated by seasonal variation.

Figure 3 plots the 7-day DFTI moving average, decomposed into three flight time components: time-at-origin (departure delay plus taxi-out time), airborne time, and taxi-in time. While airborne time is clearly the largest component of DFTI, it is evident that "spikes" in the DFTI are normally accompanied by similar "spikes" in time-at-origin, suggesting that the latter is the largest source of DFTI variation. Variation in the time-at-origin for LAX-bound flights may result from congestion at the origin airport, or the backward propagation of delays downstream, either en route or at the LAX terminal area.

Figure 1. Daily DFTI Values for LAX

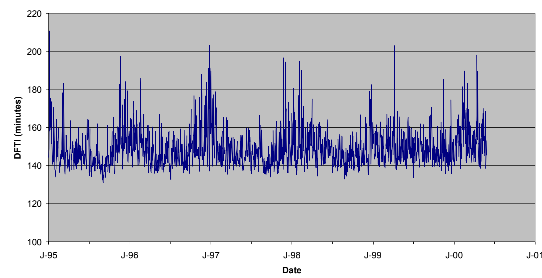

Figure 2. 30-Day Moving Average of DFTi, by Season

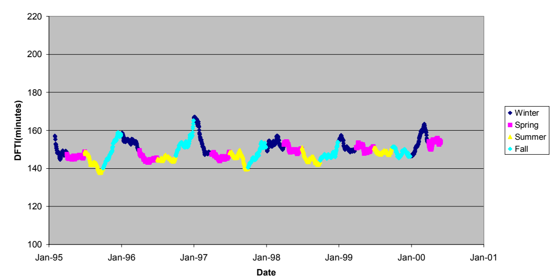

Figure 3. 7-day DFTI Moving Average, Three Components

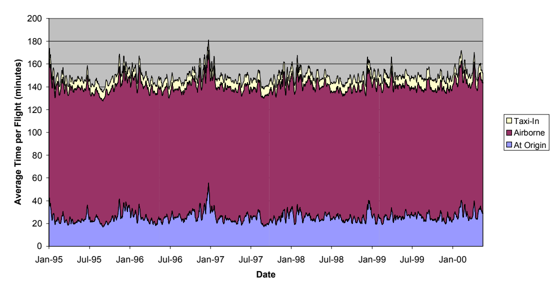

3. Weather Normalization

As with any airport, weather conditions at LAX significantly affect its capacity and thus delays for LAX-bound flights. We begin with a qualitative discussion of these effects as they occur at LAX, and then present our statistical methodology for weather normalization.

Winds at LAX are generally westerly, and as a result LAX operates in west flow 90-95% of the time. Operations change to east flow when the wind speed and angle do not allow for west flow operations. During the day west flow is in effect unless the winds preclude it. During the night (midnight till 6.30 am) LAX is running operations in east flow due to a noise abatement procedure. The capacity goes down when the airport is in east flow, because of the lack of familiarity of pilots and controllers with this configuration, and also because there are no high-speed runway exits in the east flow direction.

Winds can affect capacity even when the airport remains in west flow. If the tail wind component is strong, pilots must land at higher approach speeds, so air traffic controllers have to separate them more than usual.

The main effect of ceiling and visibility at LAX is on whether the in-board runways can be used for arrivals. This is permitted under VFR, and when departure traffic is sufficiently light. Arrival capacity increases by about 10 per hour when the in-boards are available. Separations between arriving aircraft are also somewhat less under VFR because pilots can accept separation responsibility on final approach. Operations at LAX are generally VFR if the cloud ceiling is higher than 5000 ft. If the cloud ceiling goes below 4-5000 ft, operations are IFR, but some aircraft still can get permission for a visual final approach. With cloud ceiling lower than 2-3000 ft operations are IFR only. In addition to clouds, haze—elevated polluted air—affects visibility at LAX, sometimes forcing IFR operations even in the absence of clouds. Capacity may also be reduced when the cloud ceiling is between 8000-12000 ft west of the airport, along the approach routes. These conditions compromise the ability of controllers to efficiently sequence traffic further out from the airport.

Precipitation also affects air traffic. In addition to reducing visibility, runway occupancy time increases with precipitation because it takes longer for aircraft to slow down. Both effects lead to reduced runway capacity.

As the previous discussion shows, weather conditions in the LAX area can affect our DFTI metric in a variety of ways. It is therefore important to take into account a wide range of meteorological parameters. Our task is to capture the effects of weather in a manner consistent with these complexities, and but economical enough to allow for meaningful statistical analysis.

To meet these conflicting aims we used principal components analysis, a type of factor analysis, to develop a set of daily metrics that characterize weather conditions at LAX. The data underlying these metrics was obtained from the CODAS weather database, which provides hourly data on temperature, wind, cloud ceiling, visibility, precipitation, and mode of operation (VFR versus IFR). For the factor analysis, we used the CODAS data to develop summary information for four daily periods: early morning (0-600), morning (600-1200), afternoon (1200-1800), and evening (1800-2400). For each of these periods, we calculated eight measures. These included: average temperature, visibility, and wind speed, total precipitation, proportion of time in VFR operation, and proportions of time cloud ceiling was low (under 3000 ft), medium (3000-8000 ft), and high (8000-10000 ft).

We had four daily observations for each of these eight variables, or a total of 32 variables per day. In the subsequent analysis, each variable is converted into a standardized variable (zero mean, unit variance). The factor analysis procedure was then used to collapse the 32 standardized variables to a smaller number of factors, which are also constructed so that they have zero mean and unit variance. Each of the factors is a linear combination of the original 32 variables. The factors are constructed so that the first accounts for the largest possible amount of the variation in the 32 variables, the second accounts for the largest possible amount of variation unaccounted for by the first factor, and so on. While 32 factors are needed to fully capture the variation in 32 variables, a much smaller number of factors will generally account for most of the variation. This is particularly true when the variables are highly intercorrelated, as they are in this case. Because of the intercorrelation, we were able to extract nine factors that account for 73 percent of the total variation in the original weather data set.

After extracting the factors, we rotated them in order to obtain factors that are more easy to interpret. The objective of factor rotation is to create factors that are highly correlated (either positively or negatively) with some variables while having low correlation with others. Various rotation procedures have been developed; for the weather data we chose the promax procedure. This is an oblique rotation method. Unlike the original factors, those generated by oblique rotations may be correlated with each other. Table 2 shows the correlations between the rotated factors and the original 32 weather variables. As intended, the rotated have very high correlations with certain variables and low correlations with others. Table 3 provides qualitative interpretations of all nine factors.

Table 2. Correlations Between the Rotated Factors and Original Weather Variables

VAR	TIME	FACTOR 1	2	3	4	5	6	7	8	9
Wind	Early am	-0.32	-0.06	-0.02	0.31	0.34	0.51	0.18	0.22	0.48
	Late am	-0.12	-0.06	-0.03	0.20	0.34	0.76	0.21	0.26	0.26
	Aft	0.14	0.00	0.11	0.17	-0.04	0.83	-0.04	-0.04	0.01
	Eve	-0.06	-0.12	-0.01	0.25	0.19	0.77	0.16	0.17	-0.08
Temp	Early am	0.92	-0.21	-0.18	-0.06	-0.09	0.10	-0.14	-0.13	0.01
	Late am	0.98	-0.03	-0.06	-0.02	-0.23	0.01	-0.19	-0.23	-0.04
	Aft	0.95	0.08	0.07	0.08	-0.33	-0.08	-0.24	-0.29	-0.02
	Eve	0.96	-0.03	0.00	0.02	-0.27	-0.06	-0.21	-0.22	-0.04
Visual Ops	Early am	-0.13	0.87	0.40	0.41	-0.15	-0.08	0.08	0.09	-0.12
	Late am	0.01	0.87	0.48	0.41	-0.12	-0.01	0.06	-0.22	-0.01
	Aft	0.08	0.43	0.84	0.38	-0.10	0.02	-0.00	-0.36	0.06
	Eve	-0.12	0.49	0.85	0.40	-0.01	0.05	-0.03	0.14	-0.14
Visib	Early am	0.02	0.55	0.26	0.85	0.11	0.20	0.08	0.02	-0.09
	Late am	0.01	0.53	0.37	0.88	0.11	0.21	0.12	-0.14	0.04
	Aft	0.07	0.32	0.56	0.86	0.11	0.20	0.04	-0.20	0.05
	Eve	-0.02	0.22	0.58	0.82	0.13	0.22	-0.03	0.05	-0.07
Precip	Early am	-0.15	-0.18	0.14	0.00	0.20	0.15	0.20	0.20	0.62
	Late am	-0.18	-0.12	-0.01	-0.02	0.24	0.16	0.10	0.78	0.14
	Aft	-0.14	-0.01	-0.16	-0.03	0.19	0.12	0.28	0.80	-0.03
	Ev	-0.11	0.06	-0.12	0.02	0.20	0.00	0.67	0.27	0.08
Low Ceiling	Early am	0.15	-0.87	-0.42	-0.36	0.05	0.06	-0.07	-0.11	0.13
	Late am	0.02	-0.86	-0.50	-0.36	-0.01	0.01	-0.04	0.21	0.01
	Aft	-0.04	-0.42	-0.85	-0.36	0.03	0.00	-0.02	0.32	-0.07
	Eve	0.14	-0.47	-0.84	-0.42	-0.07	-0.04	0.01	-0.16	0.12
Med Ceiling	Early am	-0.20	0.02	0.05	0.15	0.74	0.28	0.14	0.20	0.01
	Late am	-0.14	-0.08	0.02	0.16	0.86	0.22	0.14	0.12	-0.02
	Aft	-0.27	-0.19	0.01	0.11	0.78	-0.00	0.31	0.27	0.06
	Eve	-0.22	-0.15	-0.05	0.18	0.69	0.02	0.36	0.27	0.14
High Ceiling	Early am	-0.15	0.11	-0.08	0.03	0.19	0.13	0.66	0.28	-0.01
	Late am	-0.14	0.03	0.06	0.05	0.16	0.11	0.71	0.03	0.15
	Aft	-0.20	-0.03	0.12	0.18	0.20	0.09	0.71	0.08	-0.26
	Eve	-0.20	-0.09	0.16	0.12	0.26	0.17	0.24	0.16	-0.53

Table 3. LAX Weather Factor Interpretations, Rotated Factors

FACTOR	INTERPRETATION
1	Warm temperatures throughout day.
2	VFR operations and absence of low cloud ceiling in the morning.
3	VFR operations and absence of low cloud ceiling in the afternoon.
4	High visibility throughout day.
5	Medium cloud ceiling throughout day.
6	High winds throughout day.
7	High cloud ceiling throughout day.
8	Precipitation in late morning and afternoon.
9	Precipitation in early morning.

4. Demand Normalization

We use the concept of hypothetical deterministic delay (HDD) to capture the intensity of demand at LAX. The HDD concept is illustrated using a representative daily schedule into LAX, shown in Figure 4. The cumulative curve for the schedule gives the number of arrivals scheduled into LAX from the beginning of the day through the time plotted on the abscissa. The other curves represent the number of arrivals that can actually occur in light of the schedule and an assumed capacity. The total delay associated with a given schedule and capacity is the area between the scheduled curve and the appropriate constrained curve. Using this method, we calculated the daily average delay based on the schedule and a range of hypothetical capacities, varying in increments of 10 from 10 to 120. We then performed factor analysis on the 12 daily delay variables, finding that two factors, one associated with delays assuming high capacities and the other with delays assuming low capacities, capture 93% of the variation. These two demand variables, which we term "peak demand" and "base demand" respectively, are used for our normalization. Figure 5 presents examples of daily schedules with high values for peak demand and base demand, a high base demand but a low peak demand, and low base and peak demands.

Figure 4. Hypothetical Cumulative Curves

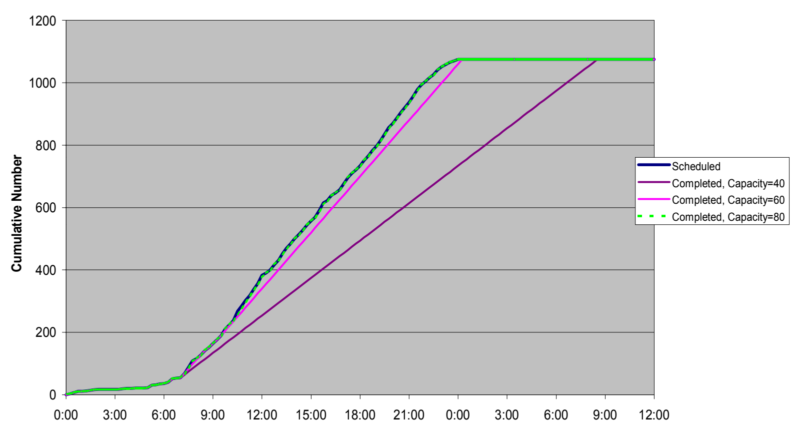

Figure 5. Cumulative Scheduled Arrivals

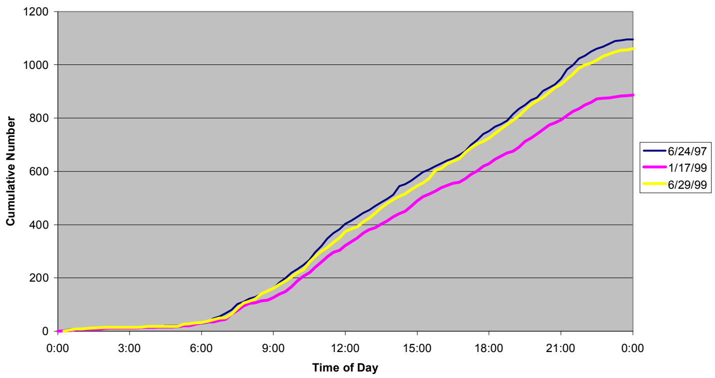

Figure 6. Origin Airport Delay Time Series for LAX

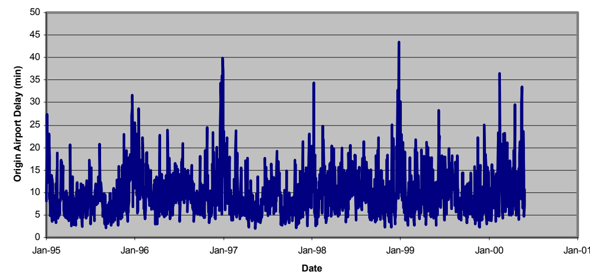

Figure 7. Origin Airport Delay 30-Day Moving Average for LAX

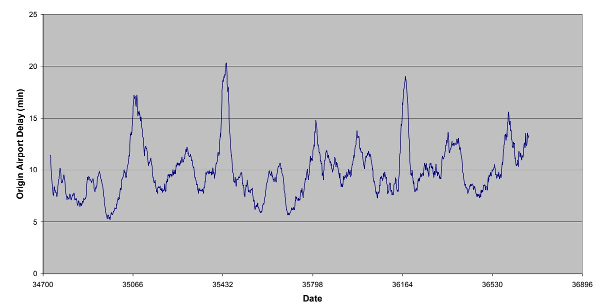

5. Origin Airport Delay Normalization

The third and final normalization variable is delay at the origins of flights bound for LAX. Obviously, if an origin airport has severe congestion that is delaying outbound flights, arrivals from that origin into LAX are likely to be effected, even when conditions at LAX are ideal. To normalize for this effect, we employ an origin departure delay index analogous to the DFTI. For every airport included in the DFTI average (see Table 1) we compute, on a daily basis, the average departure delay for all flights not bound for the LAX area (including LAX itself and all airports within 200 miles of it). Then, we compute a weighted average across all the DFTI airports, using the same weights used in the DFTI itself. In effect, this average predicts what the departure delay of an average flight to LAX would be if that flight were actually bound for another destination outside the LAX region.

Figure 6 plots the daily origin delay metric, which we label OAC for “origin airport congestion,” for LAX since beginning in January 1995. Under the best conditions, the metric is less than 5 minutes, and for the most part it remains under 15 minutes. On the worst days, it can spike as high as 30 or 40 minutes. These days are generally in the winter months, particularly January. As shown in Figure 7, which is a 30-day moving average of the same data, there is some evidence of an upward trend in this metric beginning in January 1998. After this period, the average never dips below 7.5 minutes, a value undercut fairly often in the earlier period.

6. Regression Modeling of DFTI at LAX

As explained earlier, we used the weather and demand factors as explanatory variables to model the variation of the DFTI, and its components, for LAX. As a starting point, we employed a simple linear form:

$$DFTI_t = \alpha + \sum_i \omega_i \cdot WX_{it} + \sum_j \partial_j \cdot DMD_{jt} + \theta \cdot OAC_t + \varepsilon_t$$

where:

WX_{it}	is the value of weather factor i on day t
DMD_{jt}	is the value of demand factor j on day t
OAC_t	is the measure of average departure delay for LAX origin airports
ε_t	is a stochastic error term

The weather and demand factors are the rotated ones described earlier, and the average departure delay is the variable discussed in the previous section.

We adopted the above specification as our baseline form because it is relatively simple and easy to estimate, and because estimation results are easy to interpret. As the baseline, it serves as the “hub” for other models that differ from it in various ways. We also use this specification to model the DFTI components time-at-origin, airborne time, and taxi-in time. Use of the linear form means that the coefficients in the DFTI model will be sums of the corresponding coefficients in the components models. We subsequently tried a number of alternative model specifications. These include a quadratic response surface model, a non-linear model, models in which weather is represented by more or fewer factors, and finally a non-parametric model. Here we will focus on the baseline model, with brief discussions of the alternative forms.

To complete the specification it is necessary to make assumptions about the error term, ε_t. Initial estimation results for LAX, as well as experience with similar models for other airports, reveal that the errors are heteroscedastic, with greater errors generally occurring on days when predicted DFTI values are high. The phenomenon is illustrated in Figure 8, where the predicted values for DFTI are those obtained from ordinary least squares (OLS) regression. OLS estimation on heteroscedastic data yields estimates that are unbiased, but inefficient. Moreover, the standard errors on the coefficients are biased, making it difficult to judge their statistical significance. To remedy this problem, we ran a feasible generalized least squares (FGLS) procedure. Under FGLS, we employ a prediction of squared residual as a weight on each observation. We obtained the prediction by regressing the absolute value of the OLS residuals against the explanatory variables included in (2). This estimation procedure yields results that are unbiased and asymptotically efficient.

Figure 8. Observed and Predicted Values for DFTI

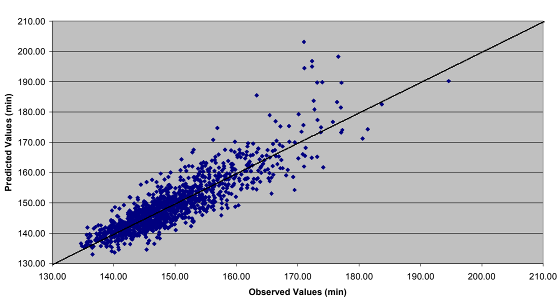

The estimation data set includes daily observations from the beginning of 1997--when CODAS became available--through May 2000. Table 4 contains the estimation results from the FGLS procedure. Every coefficient estimate but one is significant at the 0.05 level, and all but two at the 0.001 level. The adjusted R^2 of 0.74 implies that the model explains about three quarters of the DFTI variation occurring in the data set. As just explained, the accuracy of the model predictions varies, but overall the standard error of a prediction is under 5 minutes.

Table 4. Estimation Results From the FGLS Procedure, DFTI Values

Variable	Description	Estimate	T - statistic	P - value
	Intercept	138.055	567.065	0.0001
OAC	Origin airport departure delay	1.128	44.351	0.0001
WX_1	Warm daily temperatures	-1.357	-12.101	0.0001
WX_2	VFR ops, no low cloud ceiling in the morning	-0.988	-7.116	0.0001
WX_3	VFR ops, no low cloud ceiling in the afternoon	-1.123	-7.583	0.0001
WX_4	High visibility throughout day	-0.449	-3.575	0.0004
WX_5	Medium cloud ceiling throughout day	1.440	10.555	0.0001
WX_6	High winds throughout the day	0.512	4.531	0.0001
WX_7	High cloud ceiling throughout day	0.911	4.172	0.0001
WX_8	Precipitation in late morning and afternoon	1.871	8.324	0.0001
WX_9	Precipitation in early morning	-0.379	-2.614	0.0091
DMD_1	Peak demand	0.075	0.725	0.4685
DMD_2	Base demand	0.440	4.574	0.0001
ADJUSTED R^2			0.743	

We now consider the performance effects of weather, demand, and origin airport delay as they are revealed in this model. From the coefficient on *OAC* we learn that an additional minute in the expected origin departure delay adds about 1.1 minutes to the expected DFTI at LAX. This coefficient is certainly of reasonable magnitude—if the expected origin delay increases by a certain amount, one would expect the DFTI to follow suit.

All of the weather factors are statistically significant at the 0.01 level or better. Four of them, Factors 5, 6, 7, and 8, are positively related to DFTI. Since the factors are standardized (mean=0, standard deviation=1) variables, a coefficient can be interpreted as the effect of a 1 standard deviation increase in its associated variable on DFTI. Factor 8, midday precipitation, has the strongest effect. In addition to causing wet runways and lower deceleration rates, precipitation is probably an indicator of adverse wind and visibility conditions that are not fully captured by the other weather factors. Factor 5, medium cloud ceiling, also has a strong positive effect on DFTI, with a coefficient of 1.4. The influence of this factor probably derives from its impact on the ability to use the in-board runways for arrivals. Factors 7 and 6, high cloud ceiling and high winds, have weaker effects, with coefficients of 0.9 and 0.5 respectively. The latter result is certainly to be expected. High winds may force the airport to operate in east flow or cause controllers to keep larger separations in order to assure that minimums are not violated. It is less obvious why high cloud ceiling would have an impact. Generally, we may presume that the presence of such a ceiling is associated with some visibility condition, such as haze, that is not reflected in the other weather variables.

The other five weather factors are negatively associated with DFTI. Of these, the strongest impact belongs to Factor 1, temperature, for which an increase of 1 standard deviation causes a 1.4-minute DFTI reduction. The strength of this effect is somewhat puzzling. High temperature may indicate "severe clear" weather conditions with the earth's surface in full sunlight. Alternatively, the impacts may derive from effects of temperature on aerodynamic performance, since the speed of sound, and thus the speed for a particulate Mach number, increases with temperature. The other factors with negative coefficients include Factor 3, VFR operations and absence of low cloud ceiling in afternoon, Factor 2, VFR operations and absence of cloud ceiling in morning, Factor 4, high visibility throughout day, and Factor 9, precipitation in early morning. Explanations for the first three of these are obvious. As to the last, it is hard to imagine how early morning precipitation could have any direct, negative effect on flight times. The most likely explanation is that this effect occurs in the context of Factor 8, high daytime precipitation. If Factor 8 is high, then there is likely to be high precipitation in the period from 6 am to noon. If Factor 9 is also high, then this precipitation is more likely occurring during the early part of this period, when it has less impact on operations.

Estimates for the demand factors are both positive, but only the base demand factor (correlated with low capacity HDDs) is statistically significant. This implies that the DFTI is affected by the volume of operations throughout the day, not just during periods of peak demand. This is not surprising in light of the rather even temporal pattern of demand at LAX.

The above results reveal the sensitivity of the DFTI to the various explanatory variables. It is also interesting to compare the contributions of the explanators to DFTI variation. These contributions depend not just on sensitivity, but also on the degree of variation in the explanatory variable. To assess them we compute standardized regression coefficients, which relate the change in the dependent variable to the change in independent variables when both are measured in terms of standard deviations. That is, a standardized coefficient of 0.5 implies that a 1 standard deviation change in the independent variable leads to a 0.5 standard deviation change in the dependent variable. The standardized coefficients for the FGLS model are plotted in Figure 9. The standardized coefficient for origin airport delay is by far the largest, with a value over 0.6. The magnitudes of the weather factor coefficients vary between 0.04 and 0.2. The demand variables have been the least important sources of DFTI variation over the analysis period.

Figure 9. Standardized Coefficients for the FGLS Model

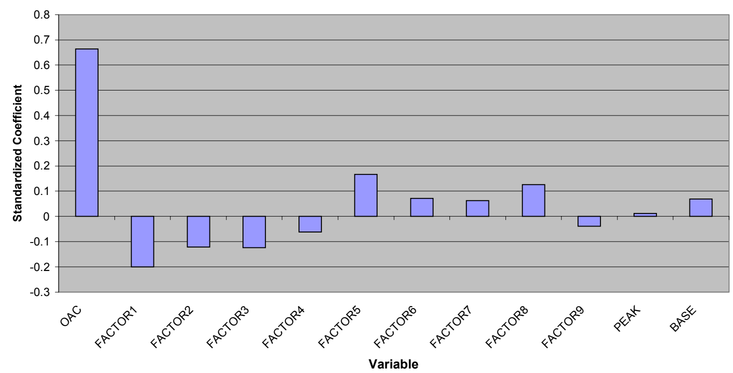

7. Regression Modeling of DFTI Components

As discussed earlier, the DFTI can be decomposed into four components: departure delay, taxi-out time, airborne time, and taxi-in time. It is therefore possible to similarly decompose the regression results obtained above. In other words, a regression coefficient in the DFTI model is the sum of regression coefficients for identically specified models of the DFTI components. By estimating the component models, we obtain additional insight concerning how the various explanatory variables influence performance.

Table 6 summarizes the regression results for the DFTI components. The three components considered are time-at-origin (departure delay plus taxi-out time), airborne time, and taxi-in time. The models were estimated using the weights used in the total DFTI model. This may not provide the most efficient estimates, since the different components may have different patterns of heteroscedasticity, but it preserves the identity between the DFTI regression coefficients and the sums of the coefficients for the DFTI components.

Table 6. Estimation Results From the FGLS Procedure, Values for DFTI Components

Variable	Time-at-origin Est	P - value	Airborne time Est	P - value	Taxi-in time Est	P - value
INTCPT	14.588	0.0001	115.6	0.0001	7.874	0.0001
OAC	1.099	0.0001	-0.012	0.4621	0.041	0.0001
WX_1	-0.065	0.4011	-1.474	0.0001	0.182	0.0001
WX_2	-0.722	0.0001	-0.233	0.0100	-0.033	0.2290
WX_3	-0.669	0.0001	-0.348	0.0003	-0.105	0.0003
WX_4	-0.201	0.0198	-0.186	0.0232	-0.062	0.0125
WX_5	0.599	0.0001	0.846	0.0001	-0.005	0.8567
WX_6	0.154	0.0480	0.428	0.0001	-0.069	0.0021
WX_7	0.372	0.0132	0.503	0.0004	0.036	0.3995
WX_8	0.897	0.0001	0.796	0.0001	0.179	0.0001
WX_9	-0.060	0.5485	-0.316	0.0008	-0.003	0.9158
DMD_1	0.034	0.6366	0.234	0.0005	-0.193	0.0001
DMD_2	0.260	0.0001	0.060	0.3367	0.120	0.0001
ADJUSTED R^2	0.804		0.427		0.213	

Not surprisingly, the time-at-origin component is most strongly influenced by average origin departure delay. In addition, seven of the nine weather factors are significant at the 5 percent level - five of these at the 1 percent level. Factors 8 and 5, precipitation and medium cloud cover, are the most important positive correlates with time-at-origin. Factors 2 and 3, morning and afternoon VFR conditions and absence of low cloud cover, are the factors that decrease time-at-origin. It should be emphasized that these weather conditions pertain to LAX, not the origin airport. Presumably, air traffic management procedures create the linkage between LAX weather and time-at-origin, through ground holds, ground stops, and other actions. This is also the case for the demand variables, of which only the base demand is statistically significant.

All nine weather factors have statistically significant impacts on airborne time. Factors 8 and 5 again have the largest positive coefficients, with Factors 6 and 7, high cloud ceiling and high winds, close behind. High temperature has by far the largest negative coefficient. Since this factor does not affect time-at-origin, it appears that the mechanism involved is not one that is governed by air traffic management actions. One possibility is that surface temperatures at LAX correlate with the upper air temperatures, which in turn affect the speed of sound, and hence the airspeed equivalent of a given Mach number. Further research is required to determine whether this or some other mechanism is at work. The morning and afternoon visibility factors (2 and 3) and early morning precipitation are also associated with lower airborne times. Of the two demand factors, only the one associated with peak demand levels is significant. This contrasts with the result for DFTI as a whole. As will be explained below, this is because the positive (and intuitively reasonable) effect of the peak demand on airborne time is offset by its negative impact on taxi-in time. Finally, and as expected, average origin departure delay is found to have essentially no effect on airborne time.

The third component, taxi-in time, does not vary much with the weather factors, although some effects are statistically significant. The temperature and precipitation factors (1 and 8) are both positively related to taxi-in time. While the latter effect is understandable, the former is more curious. One possible interpretation is that high temperatures are associated with high landing speeds (since aircraft stall speeds depend on air density), which in turn increase runway occupancy times. The good visibility factors (2 and 3) also are negatively related to taxi-in times, as is the high winds factor (6). High winds may reduce taxi-in times by allowing aircraft to land at lower ground speeds and thus exit the runways sooner. Among the demand factors, peak demand is negatively associated with taxi-in times while for base demand the association is positive. The latter may be because, when peak demands are high, there is more pressure to use the in-board runways for arrivals, and flights arriving on the in-boards are closer to the terminal and do not need to cross a runway on taxi-in. Finally, delay at the origin airport has a small but statistically significant effect on taxi-in time.

8. Alternative Model Specifications

We estimated a variety of other models for DFTI. A response surface model added quadratic and interaction terms to the linear model. Five of the 12 quadratic terms and 15 of the 66 interaction terms were found to be significant at the 0.05 level. The R^2 increased from 0.76 to 0.82. A non-linear model, in which explanatory variables were introduced in exponential form, had a slightly better fit than the baseline model. Models with different numbers of weather factors, from three to 12, were also investigated. Surprisingly, all these models, including the one with just three factors, have R^2's nearly identical to the baseline, nine-factor, model. Finally, when we used cluster analysis to identify groups of days with fairly similar weather and introduced cluster membership as a control for weather, results were somewhat less satisfactory, with R^2 values around 0.70.

9. Conclusion

A posteriori evaluations of a wide range of efforts to improve the aviation infrastructure require that we normalize performance for weather, demand, and conditions elsewhere in the system. We have presented a normalization approach based on statistical modeling of day-to-day variation in average flights times to an airport. Applying this method to LAX, we find that we can account roughly three-quarters of the observed flight time variation in terms of local weather conditions, demand, and average departure delays at origin airports. Of these, origin delay is the more important factor, followed by weather. While demand was also found to affect our flight time metric, it has not varied much over the study period and is therefore not a major source of flight time variation.

In addition to assessing total flight times, our methodology permits analysis of the various flight time components, such as time-at-origin, airborne time, and taxi-in time. In general, we find a complex web of causal relations, with weather and demand factors significantly affecting time-at-origin, airborne time, and taxi-in time. This reflects the extent to which delays at LAX propagate back to the origin airport, not only as the result of central flow control, but also through decentralized traffic management actions taken by TRACONs, Centers, and airline operations centers. This underscores the need to consider the entire flight when assessing the impact of terminal area and airport improvements, rather than just the terminal area portion of the flight.

There remains a large portion of unexplained variation in our daily flight time metric. While fairly simple models can explain about 75 percent of this variation, the more complicated specifications attempted in our research do not yield much improvement. The problem seems to be more one of omitted variables than of miss-specification Missing factors include en route winds and weather, airline operational variables, and infrastructure variables such as runway closures and navaid outages.

Biographies

Mark Hansen is an Associate Professor of Civil and Environmental Engineering at UC Berkley. He has a PhD in Engineering Science and a Masters' in City Planning from Berkeley. He obtained his undergraduate degree from Yale University in 1980, majoring in Physics and Philosophy. His research focuses on the economics and management of aviation infrastructure.

Tatjana Bolic was born in Zadar, Croatia on 1974. She has a Bachelor's degree in Transportation Engineering from the University of Belgrade in 1999, and is presently a doctoral student in transportation engineering at UC Berkeley. Her research interests include air traffic control, air traffic management, human factors.