EXAMINING THE TEMPORAL EVOLUTION OF PROPAGATED DELAYS AT INDIVIDUAL AIRPORTS: CASE STUDIES

Andrew M. Churchill, Department of Civil and Environment Engineering, University of Maryland, College Park, MD, churchil@umd.edu

David J. Lovell, Department of Civil and Environmental Engineering and Institute for Systems Research, University of Maryland, College Park, MD, lovell@eng.umd.edu

Michael O. Ball, Robert H. Smith School of Business and Institute for Systems Research, University of Maryland, College Park, MD, mball@rhsmith.umd.edu

Abstract

Delay propagation is a well-known phenomenon within the global air transportation system. Specifically, because of equipment and crew connectivity, a flight delayed early in the day can induce delays to multiple flights later in the day. In this paper we investigate this phenomenon by developing a statistical model that predicts the average flight delay after a given breakpoint time b, based on the average delay before b. We then vary b and are able to quantify the extent of delay propagation and understand its temporal evolution. We estimate this model for several U. S. airports and are able to compare these airports with respect to their delay propagation characteristics.

Introduction

At any congested airport, it is extremely important to understand how flight delays propagate over the course of the day, particularly when considering new policies or expansions that might have a significant affect on demand or capacity. For example, at some airports, a surge in demand or a capacity shortfall early in the day can be absorbed because of traditionally low usage during the early afternoon. At other airports that are more uniformly congested, early congestion can be devastating, because there is no recovery period.

It is well-known in the study of queues that for over-saturated systems, the "marginal price" of a unit of delay early in the day is much higher than it would be if that same delay were to occur later in the day. This is easily illustrated in deterministic systems, but is also generally true for (more realistic) stochastic systems.

At an airport served primarily by a small number of major carriers, such as Chicago's O'Hare International Airport, this can have implications at the individual carrier level, since each carrier is a significant contributor to the overall traffic. In these circumstances, a carrier might choose to make long-term strategic scheduling decisions such as spreading the demand over the day to avoid peaks that create congestion in the first place, or to purposefully leave vacant some "slots" in the afternoon as a buffer to allow for system recovery in the event of earlier delays. The tactical decision to cancel flights could also be affected – an early flight that is not very full, is not critical to network connectivity, and whose passengers can be otherwise accommodated, might be more likely to be cancelled if its acceptance ran a higher risk of tipping the balance of delays for the rest of the day.

At other airports with a more diverse set of servicing carriers, individual decisions such as those described above are not likely, due to the "tragedy of the commons" – it is not in anyone's best interest to unilaterally reduce their impact, as their competitors may benefit as well. The airports themselves, however, also have an interest in reducing delays, both for customer satisfaction as well as safety concerns. They might choose to adopt policies regarding the numerical values of airport acceptance rates (AARs) over the course of the day similar to the scheduling patterns described above. In particular, at high density airports where some form of slot control exists or is planned, these policies might be explicitly incorporated into the slot offering and exchange mechanism.

It is the intent of this paper, therefore, to take a careful look at empirical delay data from a variety of airports to try to understand some aspects of delay propagation at an individual airport. The central question is the following: at time t, given that some accrued delay equal to d has already occurred, what is the expected pattern of delays for the remainder of the delay? Such a question can be made more specific – for example, one might determine, for a given airport, what amount of delay represents the threshold at which a risk-based policy for congestion management might be invoked, for example a congestion-induced ground delay program (GDP).

We hypothesize that airports differ in important ways with regard to their delay propagation patterns, due perhaps to structural differences in their configurations, in the network traffic that they serve, in prevailing weather conditions that routinely affect capacity, etc. Time and distance factors can also be important – for example, some airports, like Los Angeles International Airport, serve large volumes of trans-oceanic traffic. Some airports, like Washington DC's Reagan National Airport and New York's LaGuardia Airport, are governed by explicit restrictions on the geographic distribution of the flights they serve.

The nature of our study, therefore, is largely empirical. We do not develop prediction models, but rather employ simple statistical models as tools for data mining. Our choices of models (i.e. data extraction and processing algorithms) are defended on two basic principles: a) they are simple, and b) they produce interesting results that highlight important characteristics at the airports in question, and that can be convincingly related to the causal factors described above.

In the next section, we describe how our study fits in the context of the larger body of literature on delay propagation at airports. Following that are two sections on data preparation and modeling, respectively. We then present a series of case studies at individual airports, and the paper ends with conclusions and recommendations for additional investigations.

Background

In this paper, we describe an approach to analyzing the relationship between delays earlier in the day with those that occur later in the day. Obviously, the nature of this relationship will change as the day progresses. The concept of delay propagation is very general and can be considered in many different scopes and frameworks. The methodology most appropriate for such an analysis is strongly dependent on the scope of the analysis being undertaken.

An important distinction to make is between models that analyze historical data to identify information about delay propagation that occurred, and those models that attempt to predict changes in delay propagation based on some information. The two can have markedly different applications, and, in particular, data from the first type can be used to improve models of the second type.

The approach taken in this paper is to fix a time threshold, and, given the arrival delay that accumulated before that time, examine the conditional distribution of arrival delay for the remainder of that day, using a year's worth of data. By varying the time threshold, one can start to pinpoint what parts of the day are critical in the evolution of delays, although we expect this to vary significantly by airport. This approach is fairly general, and can be used to examine the ability of an individual airport to recover from irregular operations. It includes, intrinsically, effects relating to the airport's physical configuration and climactic conditions. In addition, the scheduling and operational practices of all the carriers using that airport are incorporated. Thus, when comparing analyses from different years, if the significant carriers have changed their operations, the results of the analysis will also change.

A different approach is to examine the propagation of individual delays throughout a larger network of airports. This exercise can be useful in predicting the effects of a single delay on later operations at connected airports. Because of the size and complexity of such a model, different statistical and computational techniques are needed for analysis [1][2]. A model like this takes the stance that delay propagation between airports is inherently a system problem, and as such requires network oriented statistical techniques. Because of its structure, it is primarily oriented at capturing interactions between individual airports.

A third potential approach to this sort of analysis is to examine the propagation of delays for an individual carrier. Because an individual carrier's aircraft are likely scheduled to be used heavily all day long, an early delay should have a greater impact than a later one. However, the strength of the impact can vary strongly with the carrier's network structure and scheduling practices. A detailed analysis of this nature, using real information about route structure and connectivity of operating resources was conducted at American Airlines [3] using proprietary software. This study identified a "delay multiplier," which the authors defined as the relative value of earlier delays, in terms of induced later delay, as a function of length of initial delay and time of occurrence. A similar result (using a different methodology) will be identified in this paper for an airport-level, rather than carrier-level, delay propagation effect.

A different approach that has been undertaken by some researchers is to analyze data about individual flights recursively to identify controllable and non-controllable factors on delay propagation [4]. This study traced individual aircraft on all their flights through the National Airspace System (NAS), and removed the effects of earlier delays on later flights. Thus, the amounts and causes of individual delays could be aggregated across many flights. As it relates to this study, the amount of delay propagated from one flight to another on an individual aircraft could be readily identified, and, using this data, information about the patterns of delay propagation could be identified.

Considerable work has been done recently on the nature and effects of delay propagation, as it relates to the NAS, and individual carriers, in particular. However, there still remains room for considerable analysis of the unique nature of delay propagation effects at individual airports. In the next section, we will describe the data and the careful preparation used in this analysis.

Data Information

As described, this analysis uses historical data to examine the relationship between delay before and after some breakpoint in time. While there are many possible ways to examine the "amount" of delay that took place before or after the breakpoint, we have chosen to examine the average arrival delay per flight at a given airport. An alternative might be cumulative minutes of delay (i.e., total delay) but this quantity is confounded with the number of flights represented in the sum and the amount of time over which it was aggregated. These effects are particularly egregious at the beginning and end of a day.

At some airport *a*, a given day *d* may be partitioned into any number of time periods. We index these time periods by $t \in \{1,...,T\}$, and assume that they are of equal length. Then, for each time period *t*, many statistics about airport operations can be computed by binning individual flight records by their arrival times. The count of flights that arrived during time period *t* on day *d* at airport *a* is defined as $C_{t,d,a}$.

Now, consider arrival delays. In this paper, we will consider arrival delays relative to the schedules that airlines publish in the Official Airline Guide (OAG). Then, the average arrival delay per flight $\bar{D}_{t,d,a}$ during some time period *t* is the sum, over all flights arriving during that period, of the differences between the actual arrival time (during time period *t*) and the scheduled arrival time (not necessarily during time period *t*), divided by the total number of flights which arrived during that time period, $C_{t,d,a}$. Basically, when assigning the delay of an individual flight to a time bin, one has to choose whether to assign it according to either its scheduled arrival time or its actual arrival time, and in this case we have chosen the latter, mostly because this choice is most consistent with the standard sources of data for flight delays.

One small technicality has to be dealt with concerning congested periods that traverse midnight, and hence overlay two consecutive days. Since, in this paper, we are trying to examine the relationship between delay earlier in a day and delay later in a day, we must exclude delay that happens shortly after midnight. That delay, while technically occurring on day *d*, is actually related to, and should be considered with, day $d-1$. As such, a day, in the context of this paper, refers to the time period between 0400 local time on day *d* and 0400 local time on day $d+1$. The index *d* refers to the calendar day during which the majority of the time periods take place. As an example, in the context of this paper, the data for January 2 would exclude data between 0000-0400 local time on January 2, but would include data from 0000-0400 local time on January 3. However, January 2 would still be indexed as $d = 2$, despite the inclusion of data from the subsequent day.

The data for this paper was drawn from the Aviation System Performance Metrics (ASPM) database, which is maintained by the Federal Aviation Administration (FAA). The quarter hour Airport Analysis databases were used. If for any reason the reliability of this aggregate data were to be questioned or if a different set of reporting airports were required, individual flight records could be used to reproduce the data.

Because we are comparing delay data aggregated before a given time breakpoint with similar data aggregated afterward, we must carefully define which data are included in our definitions of prior and post breakpoint time in order to facilitate a meaningful analysis. The prior delay $\bar{D}_{prior,d,b_i,a}$ is the average delay experienced by each flight that arrived to airport *a* between the beginning of the day (0400 local) and the end of time period *b* on day *d*. That is, given time-varying sequences of average arrival delays $\{\bar{D}_{t,d,a}\}$ and airport arrival counts $\{C_{t,d,a}\}$ over some day *d*, the prior delay is calculated as the weighted mean of average arrival delays, as shown in (1).

$$\bar{D}_{prior,d,b,a} = \frac{\sum_{t=1}^{b} \bar{D}_{t,d,a} C_{t,d,a}}{\sum_{t=1}^{b} C_{t,d,a}} \qquad (1)$$

The post delay $\bar{D}_{post,d,b,a}$ is defined in a manner similar to prior delay: the average delay experienced by each flight which arrived during time period $b+1$ or greater, but before the final time period *T* (which ends at 0400 local), of day *d*. This is shown in (2).

$$\bar{D}_{post,d,b,a} = \frac{\sum_{t=b+1}^{T} \bar{D}_{t,d,a} C_{t,d,a}}{\sum_{t=b+1}^{T} C_{t,d,a}} \qquad (2)$$

The prior and post arrival delays are calculated for each possible combination of breakpoint time *b* and day *d*, for a given airport *a* over some longer time period (for example, the course of one year). Until this point, the procedure described thus far is simply an aggregation of historical data.

As an example, consider January 7, 2005 at Los Angeles International Airport (LAX). Shown in Figure 1 is the information described previously, for the case $b = 10$. The bars are the average arrival delays accounted for in each quarter hour. The darker bars are those assigned before the breakpoint time (1000 local), and the lighter ones are those that occur afterwards. The black curve is the quarter-hourly count of arrivals, which is used here as the weighting function for the delays.

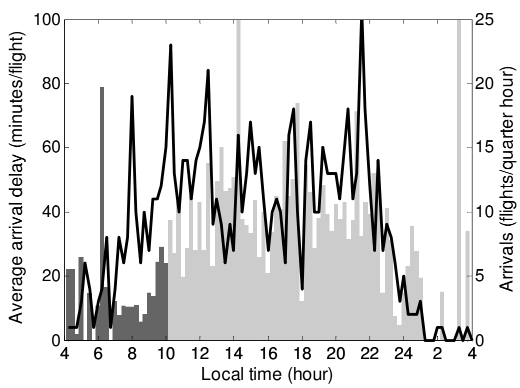

Figure 1 - Sample arrival delay/count information

The information shown in Figure 1 is processed as described and the result, for the selected day (January 7, 2005), breakpoint (1000 local), and airport (LAX), is shown in Figure 2.

While it may seem that a significant amount of fidelity has been lost in the procedure (i.e., the transition from Figure 1 to Figure 2), several things must be remembered. First, this procedure will be repeated for each possible breakpoint time *b* on a given day. As such, some semblance of the profile shape of the delay and arrival count throughout the day is retained. Second, this is not an effort focused on a micro-scale analysis of every variation in flight delay. We are seeking to identify and explain trends and patterns only in the propagation of delay at a single airport.

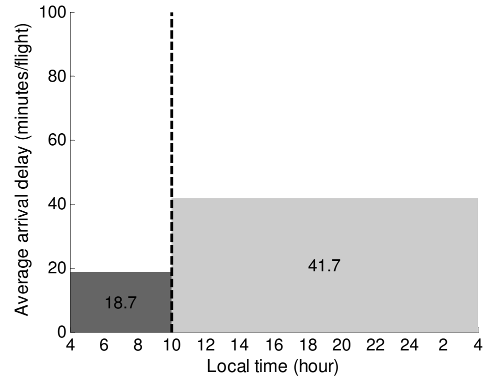

Figure 2 - Prior and post breakpoint time arrival delays for *b* = 10

Modeling Approach

Figure 3 shows a scatterplot of 365 data points for the prior and post delays at San Francisco International Airport for the year 2005, with a breakpoint time of $b = 10$. The horizontal axis shows the average prior delay, while the vertical axis shows the average post delay. For a different breakpoint time, later in the day, one would expect a distribution of prior delays that included larger values, and perhaps the extreme values of the post delay would become correspondingly smaller. For each value of the prior delay, one could imagine taking a vertical slice through the figure, and the set of post delays represents the conditional distribution of post delay, given the value of the prior delay.

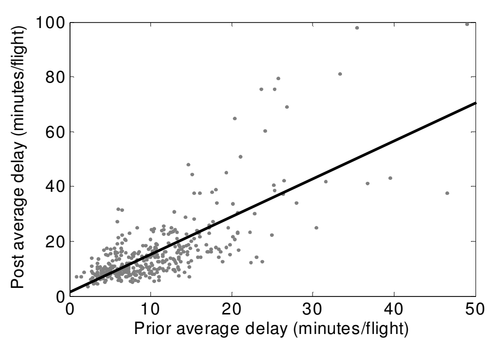

Figure 3 - Scatter plot for SFO in 2005 for *b* = 10

Importantly, Figure 3 shows a definite increasing trend in post delays as the prior delay grows. This makes perfect qualitative sense for a congested airport: the worse things are earlier, the greater the deleterious impact that will have on congestion later in the day. Of course, there comes a time when this trend is not retained, specifically at the end of the day when the demand subsides and delays decrease strictly because of the scarcity of demand, as will be observed later.

We use a simple linear regression model to capture this increasing trend. More complicated models could be proposed, but for now, very interesting results seem to come from even this simple choice. More specifically, for a fixed breakpoint time *b* and airport *a*, with the post and prior delays varying between days, we estimate the linear regression model shown in (3) using the ordinary least squares (OLS) method.

$$\bar{D}_{post,d,b,a} = \alpha_{b,a} + \beta_{b,a}\bar{D}_{prior,d,b,a} + \varepsilon_{d,b,a} \qquad (3)$$

The index for observations is the day *d*. The parameter $\alpha_{b,a}$ will be referred to as the model intercept, and this can be interpreted as some form of "background delay," since it represents delays whose existence and magnitude is independent of the changes in prior delay. The coefficient $\beta_{b,a}$ is referred to as the model slope, and it is very similar to the "delay multiplier" described in the American Airlines study [3], in the sense that it represents the expected marginal contribution to later average delays by a unit increase in average delay earlier in the day. Each of these model parameters is conditioned on a given breakpoint time period b and airport *a*. Typical assumptions are used about the error term $\varepsilon_{d,b,a}$.

This model is estimated for each possible breakpoint time *b* at a given airport *a*. The resulting model intercepts and slopes create sequences of slope and intercept values: $\{\alpha_{b,a}\}$ and $\{\beta_{b,a}\}$, respectively. Another possible sequence of interest for examination after estimating the model for each breakpoint time is that of the coefficient of determination, R^2. As this is a single variable regression model, the resulting sequence $\{R^2_{b,a}\}$ can be used to examine the strength of the correlation between the dependent and independent variables. Each of these sequences is of length *T*.

Airport Case Studies

In this section, we present results for several airports in the United States with interesting delay propagation effects, and make comparisons between the results for each.

LGA: La Guardia Airport

At New York's LaGuardia Airport, scheduled operations regularly exceed the best nominal airport acceptance rate. The precariousness of the demand to capacity ratio was made quite evident when increases in demand were allowed under the "AIR-21" legislation and extreme delays resulted [5]. As such, even in good weather conditions, serious delays are often observed. This high level of scheduled operations is shown in Figure 4.

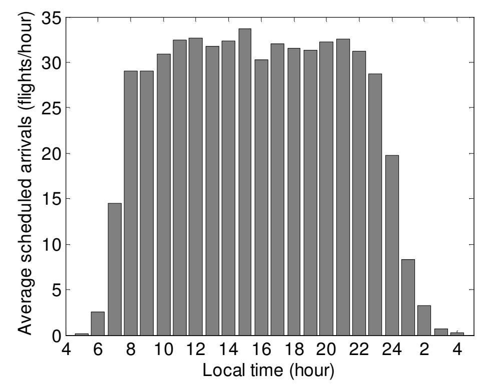

Figure 4 - Average scheduled hourly arrivals at LGA for 2005

Intuitively, one would theorize that such a uniformly high level of scheduled operations should exact a very severe penalty for early delays. In addition, because LGA is served by many carriers, and is not used as a network hub, the phenomenon described previously in which carriers may be more reluctant to make schedule adjustments may manifest itself. As shown in Figure 5, this is in fact the case. The model slope, or the cost of early arrival delays in terms of later ones, rises quickly and stays above 1.0 for much of the day. The value of 1.0 is important, since periods for which this is true exhibit non-linear impacts to later delays from earlier delays. The airport does not have a slack period in the schedule during the interior of the operating day during which to recover from these early delays. Thus, the only recovery period is late at night after the scheduled demands have been exhausted.

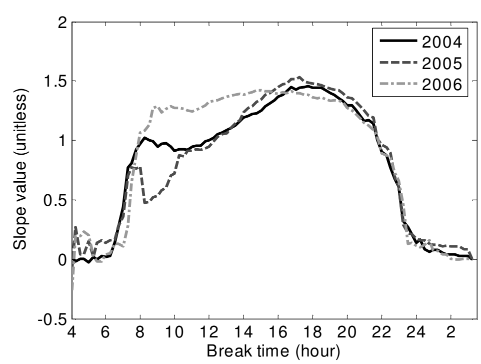

Figure 5 - LGA slope curves

Another characteristic evident in the results for LGA is a spike that occurs between 10pm and midnight in the intercept curves, as shown in Figure 6. Intuitively, this spike seems strange, since traffic at that hour at LGA is extremely limited, and the airport should be able to rapidly recover from delays. In fact, the phenomenon being observed is the fact that the airport has a scheduling curfew at the end of the day, and so, any traffic arriving after that hour is, by definition, delayed. This phenomenon will also be observed for DCA. Both airports have scheduling curfews for the obvious reason that they are both slot-controlled.

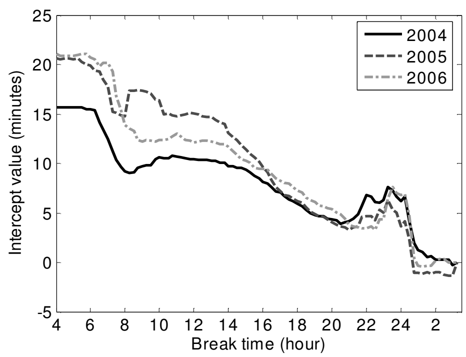

Figure 6 - LGA intercept curves

ORD: O'Hare International Airport

As with New York's LaGuardia Airport, Chicago's O'Hare International Airport is highly congested, and often experiences significant levels of delay. Like LGA, it is scheduled very fully, as shown in Figure 7.

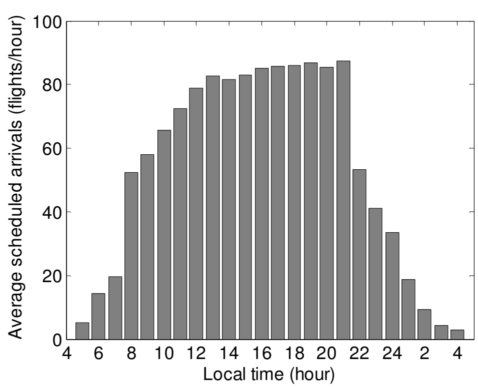

Figure 7 - Average scheduled hourly arrivals at ORD for 2005

Unlike LGA, ORD is used as a hub for two major carriers, potentially resulting in different schedule adjustment decisions during irregular operations. Because ORD is regularly scheduled at, or beyond, its nominal capacity, there exist no interior periods during which accrued delays can dissipate. This is similar to the phenomenon observed at LGA, and is shown in Figure 8. The reduction during 2006 is most likely explained by the negotiated schedule reductions that took place over the prior two years, and also the demise of Independence Air.

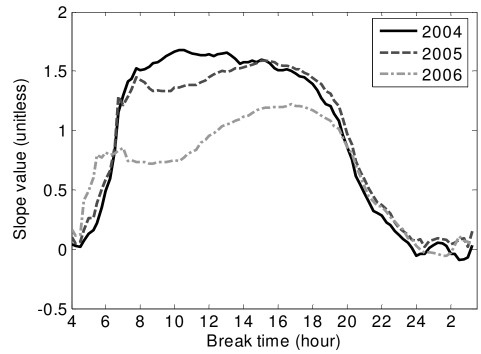

Figure 8 - ORD slope curves

It is reassuring to examine the plot of the sequence of R^2 values at ORD. Intuitively, one would theorize that this graph should rise fairly linearly, since, as the day progresses, more information about the delay pattern is emerging, and the predictions of later delays should become more and more accurate. As shown in Figure 9, this is true for this particular case study. However, as discussed later, this is not always the case in other situations. In Figure 9, as in all plots of R^2 values, the curves drop sharply at the end of the day, because the fall-off in schedule means that later delays will inevitably fall to zero, regardless of what happened earlier in the day.

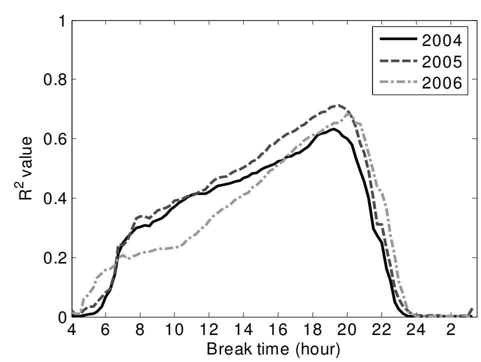

Figure 9 - ORD R^2 curves

ATL: Hartsfield-Jackson International Airport

The results for Atlanta's Hartsfield-Jackson International Airport (ATL) are remarkably similar to those for ORD. This result might be expected, since the two airports have many similarities. They are both very large, very congested, and have extensive hubbing operations for major carriers. The average schedule for 2005 is shown in Figure 10.

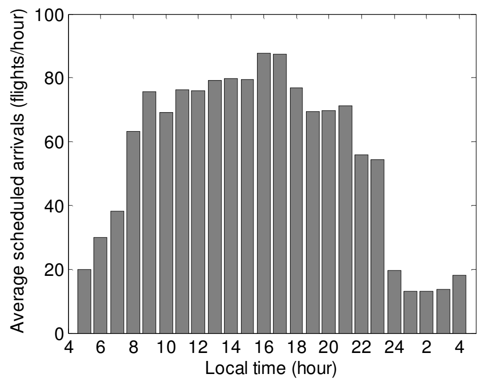

Figure 10 - Average scheduled hourly arrivals at ATL for 2005

The slope curves shown in Figure 11 for ATL look very similar to those that were shown in Figure 5 for ORD. They quickly rise above 1.0 and stay there until evening time, at which time they drop off very far. It is logically consistent that the slope curves should drop off at the end of the day, given the smaller number of operations occurring at those hours. More importantly, however, as evidenced by the R^2 plots late in the day, the regression model form simply isn't appropriate late in the day, since the driving force is the vanishing of scheduled demand.

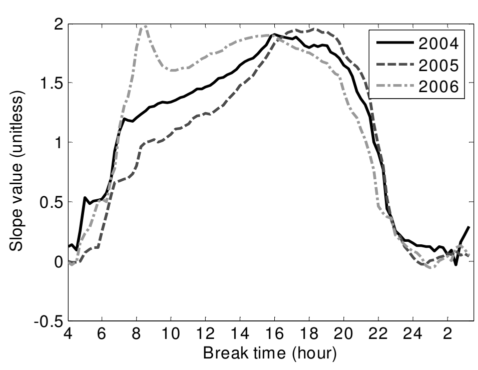

Figure 11 - ATL slope curves

The R^2 curves for ATL shown in Figure 12 look remarkably similar to those shown in Figure 9 for ORD. They tend to rise steadily and nearly linearly to a fairly high peak where the correlation between earlier and later delays is highest. This suggests that delays at large hub airports, such as ATL and ORD, are somewhat easier to predict than those for smaller airports, at least for those periods in the meaty part of the schedule when plenty of earlier flights have been able to generate delays, and sufficient later flights remain to suffer the consequences thereof. This view is consistent with the idea previously discussed that the presence of one (or a small number of) large carriers dominating an airport tends to make the operations there more predictable.

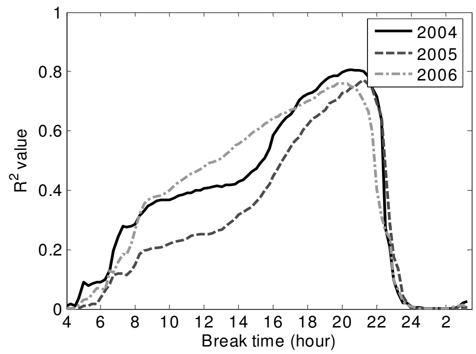

Figure 12 - ATL R^2 curves

LAX: Los Angeles International Airport

LAX is an interesting airport for analysis because it is a very busy one (as illustrated in Figure 13), but it is not a hub for any of the legacy carriers. In addition, it accommodates significant numbers of long-distance international flights, which may result in different operating paradigms.

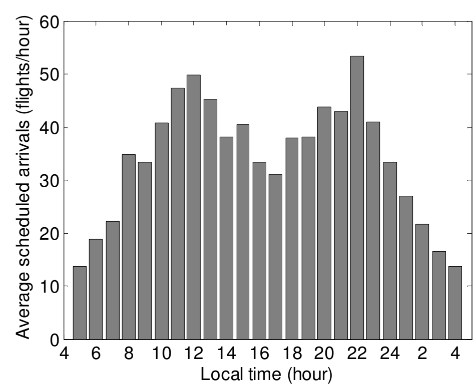

Figure 13 - Average scheduled hourly arrivals at LAX for 2005

The plots of R^2 values for the successive breakpoint models at ORD matched intuition. However, at LAX, they take a different shape. As shown in Figure 14, the curves tend to rise to a peak, and then remain flat at, or near, the same value. While this behavior is not immediately explicable, it may be important when managing irregular operations at LAX to know that one's ability to predict later delays peaks in the morning and does not markedly improve as the day progresses. In any event, the R^2 values during this period are relatively good, at least from the perspective of delay forecasting in the aviation domain.

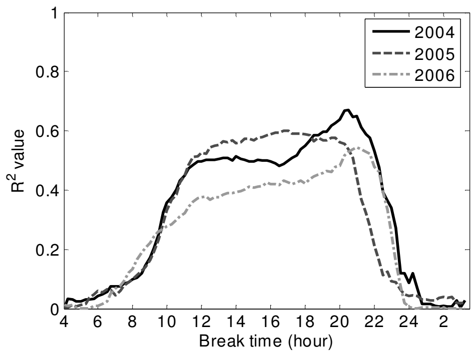

Figure 14 - LAX R^2 curves

The slope curves for LAX are also interesting to examine. They are shown in Figure 15. They tend to rise less quickly than those for the major hub airports (ORD, ATL) discussed previously. The curves still exceed 1.0 and remain there for much of the day. The likely explanation for this slower and later rise in the morning hours is that LAX does not have the same tight scheduling pressure as do the hub airports. The greater slack in the schedule prevents delays from propagating as strongly.

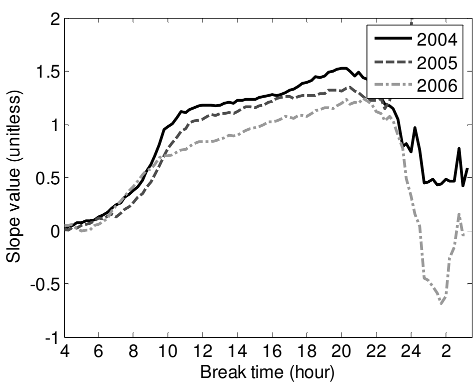

Figure 15 - LAX slope curves

DCA: Ronald Reagan Washington National Airport

Washington DC's DCA is another small and highly constrained airport, like LGA. Its average hourly schedule for 2005 is shown in Figure 16.

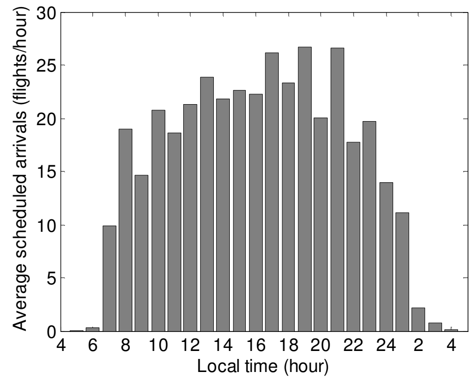

Figure 16 - Average scheduled hourly arrivals at DCA for 2005

As mentioned for LGA, airports of this type with evening curfews tend to show a spike in the intercept curves late in the evening. This phenomenon is seen again for DCA in Figure 17. Because flights arriving at these late hours are, by definition of the curfew, guaranteed to be late, the largest contribution to their delay comes not from the delays that occurred earlier in the day, but solely from the fact that they are arriving, even if only a few minutes, past the curfew time.

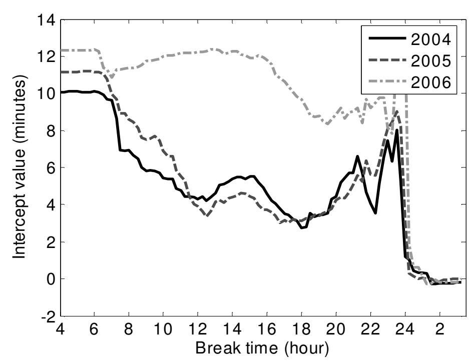

Figure 17 - DCA intercept curves

SFO: San Francisco International Airport

SFO is not scheduled as fully as some of the other airports discussed here, as can be seen in the average schedule shown in Figure 18. On the other hand, SFO regularly experiences IMC conditions and the IMC arrival capacity is roughly half of the VMC capacity.

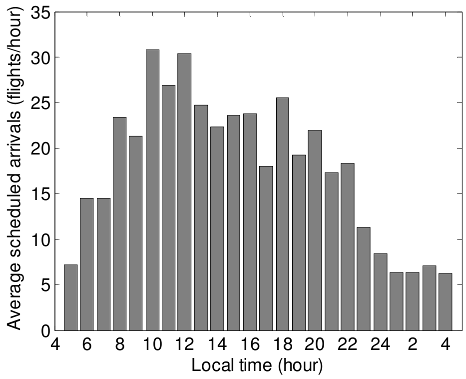

Figure 18 - Average scheduled hourly arrivals at SFO for 2005

IMC capacity is most common during the morning hours due to the presence of marine stratus conditions. As a result, morning delays resulting from severely decreased arrival capacity are extremely common. The peak in the slope curve in Figure 19 roughly corresponds to the earlier demand peak, i.e. around 1000 local time. At this time there is high demand but it is less likely that the marine stratus has burned off. While there is also high demand around 1200, at this time it is more likely that the airport is back to VMC capacity. Thus, the slope is lower.

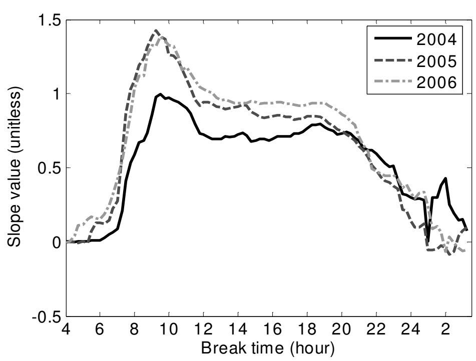

Figure 19 - SFO slope curves

Conclusions

This paper has introduced a simple statistical tool for manipulating historical delay information at individual airports to investigate various aspects of the dynamics of delay propagation over the course of a day. The intent is not to use the tool as any kind of predictive model, but rather to gain insights into the process of delay propagation, and in particular to compare across different airports, with very different operating paradigms, to see what structural effects from those airports manifest themselves in the observed patterns of delay propagation.

Three primary statistics can be plotted over the course of a day for an airport: the intercept coefficient, which represents in some sense the quantity of delays that are caused by reasons other than propagation of earlier delays; the slope coefficient, which represents the marginal rate at which later average delays result from earlier delays, and the R^2 goodness-of-fit-statistic, which represents the quality of the relation between earlier and later delays.

The slope coefficient is very similar to, but smaller than, the American Airlines notion of "delay multiplier." Since the slope coefficient is applied at the airport level and the delay multiplier for an individual carrier, it is likely that the slope coefficient is smaller since airports (even overscheduled ones) have more slack than do carrier schedules. Airports have degrees of freedom, particularly when they are served by multiple carriers, but carrier schedules are strongly constrained by the concerns of connecting flights and maximizing the utilization of aircraft.

The R^2 profiles uniformly decline sharply at the end of the day, simply because exhaustion of the schedule implies that accrued delays will inevitably be absorbed, regardless of their magnitude. This decline can be observed very sharply at slot-controlled airports or other airports with scheduling curfews. In any event, once this trend has begun, the values of the slope and intercept coefficients for corresponding time periods at the same airport are obviously unreliable, but this is nearly moot since delay propagation is not a concern at the end of the day.

This paper contributes a number of small insights into the temporal dynamics of delay propagation at airports. Importantly, the methods are simple, but nonetheless produce outputs that are rich in detail and whose patterns can be tied explicitly to important aspects of the operating paradigms at the airports. This is an important contribution because it helps understand the role of these policy decisions in the evolution of delays at an airport.

Another contribution is the fact that the R^2 profiles show the strength of correlations between early and late delays across the course of the day. While this paper is not concerned with estimating delays and their propagation, certainly many other papers are, and this helps provide insight as to when and where those efforts are most likely to be of high statistical quality.

Finally, a small point is that the plots reinforce the fact that early delays have a higher marginal deleterious effect on later performance at the airport. This has been cited in numerous places before, and is strongly supported by theory. It is important to be reminded of this fact, however, to encourage scheduling strategies and operating policies at the airport that are sensitive to the intricacies of queuing system delays, that hedge properly against periods of poor performance. It may also be politically expedient to remember this fact if it is necessary to introduce policies at congested airports, such as strict slot controls, that mitigate against this effect in ways that might seem counter-intuitive to the uninitiated, such as providing explicit slack periods for delay recovery.

Acknowledgements

This work was supported by the National Center of Excellence for Aviation Operations Research (NEXTOR), under contracts from the Federal Aviation Administration (FAA). Opinions expressed herein do not necessarily reflect those of the FAA.

References

[1] Xu, N., G. Donohue, K.B. Laskey, C.-H. Chen, 2005, “Estimation Of Delay Propagation In Aviation System Using Bayesian Network,” 6th USA/Europe Seminar on Air Traffic Management Research and Development, Baltimore, Maryland.

[2] Laskey, K.B, N. Xu, C.-H. Chen, 2006, “Propagation of Delays in the National Airspace System,” 22nd Conference on Uncertainty in Artificial Intelligence, Cambridge, Massachusetts.

[3] Beatty, R., R. Hsu, L. Berry, J. Rome, 1998, “Preliminary Evaluation of Flight Delay Propagation through an Airline Schedule,” 2nd USA/Europe Seminar on Air Traffic Management Research and Development, Orlando, Florida.

[4] Wang, P., L. Schaefer, L. Wojcik, 2003, “Flight Connections And Their Impacts On Delay Propagation,” 22nd IEEE Digital Avionics Systems Conference 2003, Indiana.

[5] Fan, T. P. and A. R. Odoni 2002, “A Practical Perspective on Airport Demand Management”, *Air Traffic Control Quarterly*, 10. 3, pp. 285-306.

Keywords

Air traffic delays, delay propagation, airport operations, airport management, air traffic demand, airport scheduling

Biographies

Andrew M. Churchill is a Ph.D. student in the Department of Civil and Environment Engineering at the University of Maryland, where he serves as a Graduate Research Assistant for NEXTOR. He received his B.S. degree in Aerospace Engineering from the University of Maryland. He has previously worked in the airline industry.

David J. Lovell is an Associate Professor in the Department of Civil and Environmental Engineering at the University of Maryland. He holds a joint appointment with the Institute for Systems Research. Dr. Lovell received his B.A. degree in Mathematics from Portland State University, and his M.S. and Ph.D. degrees in Civil Engineering from the University of California, Berkeley.

Michael O. Ball is the Orkand Corporation Professor of Management Science in the Robert H. Smith School of Business at the University of Maryland. He also holds a joint appointment within the Institute for Systems Research in the Clark School of Engineering. Dr. Ball received his Ph.D. in Operations Research in 1977 from Cornell University. He is co-Director of NEXTOR, and he leads the NEXTOR Collaborative Decision Making project.