[th] USA/Europe ATM R&D Seminar 23-27 June 2003, Budapest, Hungary

SPECTRAL ANALYSIS OF AIRPORT CAPACITY AND DELAY

Jerry D. Welch and Shafique Ahmed, MIT Lincoln Laboratory, Lexington, MA

Abstract

Objective delay statistics are sensitive measures of the effect of capacity improvements at airports [1]. However, seasonal, economic, meteorological, and operational changes can mask improvements as significant as new runways [2]. A case in point is the runway that opened in October 2000 at Phoenix. Comparison of average monthly delays before and after October shows increased delay with the new runway. However, spectral analysis of airport throughput and delay is a new way of examining airport performance that separates the positive effect of the new Phoenix runway from the background noise caused by en route congestion, weather, and flow control actions.

The histogram of the number of times each hourly arrival count occurs at an airport is roughly analogous to the Fourier transform of the arrival time series. We call this histogram the occurrence spectrum for the airport. We call related distributions of delay versus throughput delay spectra. Spectral analysis is a useful discriminant because airport queuing relates more strongly to delay at the high throughput end of the spectrum. At hubs, delay relative to schedule decreases and queuing delay increases at the highest end of the spectrum because very high throughput occurs when on-schedule aircraft arrive during hubbing rushes. At non-hubbing airports, delay relative to schedule sometimes increases and queuing delay decreases at the highest end of the spectrum because very high throughput occurs when delayed flights merge at the airport with on-time flights in favorable weather. This paper analyzes the patterns of throughput and delay spectra for important US airports that operate near capacity.

Airport spectral analysis provides useful new insights into the benefits of congestion management initiatives. Spectral analysis may also help classify airports for simulation studies, help track airport performance changes over time, and help improve airline scheduling accuracy.

This work was performed for the National Aeronautics and Space Administration under Air Force Contract No. F19628-00-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by NASA.

Introduction

In this paper, we use FAA Aviation System Performance Metrics (ASPM) flight data [3] for spectral analysis of airport capacity and delay. The spectra are distributions of occurrence counts and averages of delays as functions of throughput. Spectral analysis of airport delays provides a filtering tool for separating airport delays from en route delays. Terminal queuing is the largest component of queuing delay when the arrival throughput approaches the airport capacity, whereas delay at the low throughput end of the spectrum is usually associated with en route effects. As a case study, we use airport spectral analysis to help clarify the effects of the new runway at Phoenix International Airport (PHX). Comparison of mean monthly delay before and after the runway opening has indicated increased delay after the new runway. It is possible that other confounding factors caused mean delay to increase with time. However, comparison of before and after occurrence and delay spectra provides additional information on which to base a benefit conclusion.

FAA Delay Metrics

We use FAA Aviation System Performance Metrics data for spectral analysis of airport capacity and delay. ASPM provides two distinct data sources. We refer to ASPM files of delay and operations data as ASPM/D. They count and average events in 15-minute periods for 50 major US airports. The 15-minute ASPM/D operations counts labeled "Departure Count use for Score Card Calc" agree well with departure counts derived from terminal radar surveillance data at Dallas Fort Worth Airport (DFW) [4]. ASPM also provides files of individual flight data, which we refer to as ASPM/F.

ASPM computes and reports delay averages by assigning zero delay to all early arrivals. When negative delays are lumped together with zero delays in this way, the mean arrival error increases, the variance decreases, and information on detailed delay behavior is lost. However, one can compute the full two-sided delay from the ASPM/F data for individual flights. ASPM reports several delay values. In this paper, we focus on ASPM "arrival" delay and "airborne" delay. Arrival delay is measured relative to the scheduled gate arrival time. Airborne delay is measured relative to the flight time estimated in the final amended flight plan on departure. We have shown previously that airborne delay is a good surrogate for queuing delay [5,6].

Figure 1 shows the distribution of two-sided arrival delay for all flights during the month of April 2000 into DFW from the other ASPM airports.

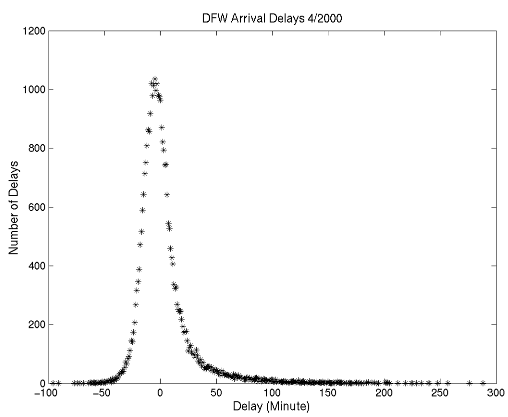

Figure 1. Distribution of two-sided ASPM arrival delay (actual gate arrival time - scheduled gate arrival time) for DFW in April 2000

Early arrivals (negative delays) were common; some flights arrived as much as an hour early, and the mode of the distribution was at negative 5 minutes. However, because more flights were very late than very early, the mean arrival time error was positive, at 2.05 minutes per arrival. On average, aircraft operators in April 2000 accurately scheduled arrivals to DFW. However, the standard deviation of the arrival time error was 23.6 minutes.

Figure 2 shows the two-sided airborne delay distribution for all flights to DFW during April 2000 from the other ASPM airports.

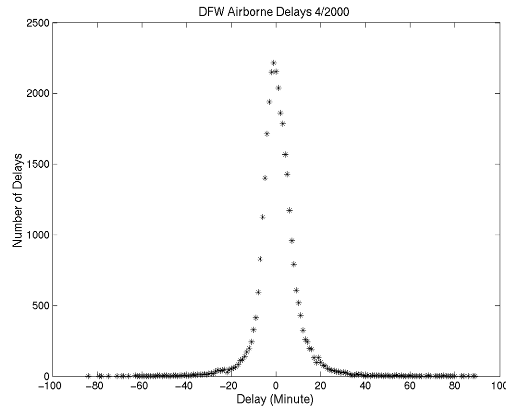

Figure 2. Distribution of two-sided ASPM airborne delay (Actual Airborne Time – Flight Plan Estimated Time En Route) for arrivals at DFW in April 2000

As with the arrival delay distribution, negative delays occurred frequently. However, the airborne delay distribution did not skew as strongly to positive delays as did the arrival delay distribution. The mean airborne delay was small, at 0.74 minutes per aircraft. On average, aircraft operators flying to DFW in April 2000 accurately predicted en route flight times at takeoff by accounting both for winds aloft and flight time history.

However, the standard deviation of the two-sided airborne delay at DFW in April of 2000 was 9.2 minutes per arrival. This is about twice the standard deviation of the one-sided airborne delay.

Airport Spectra

Here, we use the two-sided Arrival and Airborne Delay to extend an analysis introduced at ATM 2001 [5] of the relationship of delay to throughput at key airports. The purpose of that investigation was to determine if airborne delay increases with throughput as predicted by steady state queuing theory. Previously, we measured throughput by counting arrivals in 15-minute windows. We found that at Dallas-Ft. Worth (DFW) and Newark (EWR) the average one-sided airborne delay for all arrivals in a 15-minute window was essentially independent of the number of arrivals in that window. We now average the two-sided delay and the throughput over 1-hour windows to provide better correlation between the current runway throughput and the terminal delay suffered by each arrival. Figure 3 plots the relationship between two-sided delay and throughput for all arrivals to DFW during calendar year 2000.

Each plotted delay value is an average of the individual delay averages for all hourly windows that experienced the same arrival throughput. We step the hourly windows through all of the 15-minute periods in the year. As we found previously at DFW when we examined single-sided delay, arrival delay decreases with throughput, and airborne delay is essentially independent of throughput. Uniform airborne delay implies that DFW has excess runway capacity and that more flight delay occurs upstream (in en route airspace) than on approach to the airport runway. Uniform delay is consistent with the concept of flow control as an attempt to maintain a constant level of congestion in terminal airspace.

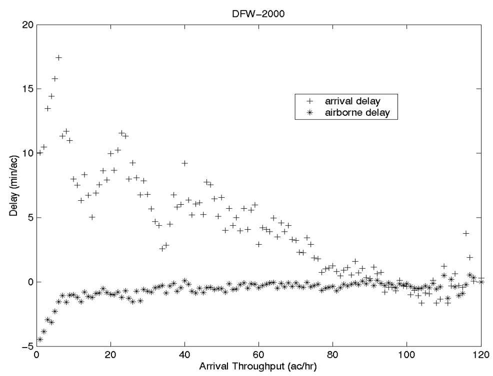

Figure 3. Annual ASPM two-sided arrival and airborne delay versus arrival throughput at DFW in CY2000.

Figure 3 shows a new effect that did not appear in our previous ASPM single-sided delay averages. On average, the airborne delay decreased significantly in those hours in which fewer than seven aircraft landed. Thus, an aircraft that arrived at DFW with little or no contending traffic was likely to experience very low airborne delay.

At higher arrival rates, the variation of two-sided delay with throughput was similar to what we reported previously for the variation of single-sided delay. As the landing rate increased, airborne delay remained almost constant, and arrival delay decreased, eventually falling to about the same level as the airborne delay in those periods when the arrival rate exceeded 90 arrivals per hour.

Figure 4 is the corresponding plot of the number of occurrences of all observed arrival rates at DFW for the year 2000. It also notes the number of stepped hourly periods with no arrivals. The peak throughput observed was 120 arrivals in a 1-hour period. As in Figure 3, hourly sliding-window counts for all 15-minute periods in the year 2000 determine the arrival rate. The occurrence distribution tells us that the annual delay statistics become statistically insignificant at arrival rates above about 115 per hour. This causes the "noise" in the average delay values above that arrival rate.

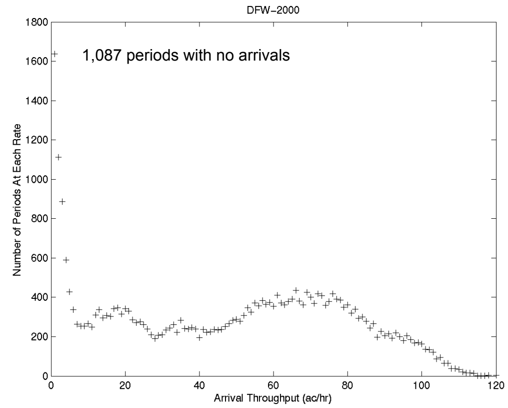

Figure 4. Frequency of occurrence of observed arrival throughputs at DFW in CY2000.

We refer to these charts as *spectra* because plotting the airport behavior as a function of arrival throughput allows us to analyze the performance of each airport as a function of arrival rate.

Spectral Analysis of Airports

Figures 5 through 19 are spectral charts for the 15 airports with the highest demand/capacity ratios in calendar year 2000. The airports appear in alphabetical order by airport identifier. Each figure includes an occurrence spectrum chart and a separate delay spectrum chart, as plotted above for DFW.

All of the airports were similar in one important way: they did not experience rapid growth in airborne delay as throughput increased. Most showed only a slight increase in airborne delay with throughput, and four (LGA, MSP, ORD, and SEA) experienced slight decreases in airborne delay at the highest end of the throughput spectrum.

As noted above in our observation on DFW, airborne delay at low throughputs seems to reflect slowdowns caused by inadequate capacity in the en route sectors feeding the airport. Flow control procedures that reduce effective en route capacity in order to relieve pressure on the airport can cause flight delay at low throughputs. Hub airports have another mechanism for maintaining constant airborne delay. Hubs routinely alternate between arrival and departure configurations. This periodic re-configuration is equivalent to an intentional variation in arrival capacity to match scheduled changes in arrival rate. It keeps airborne delay relatively independent of throughput.

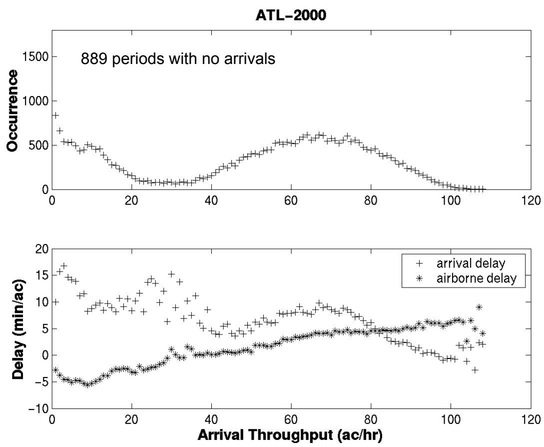

Figure 5. Spectra for Atlanta (ATL).

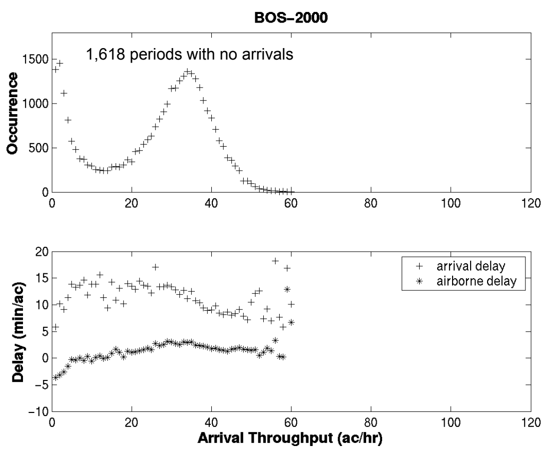

Figure 6. Spectra for Boston (BOS).

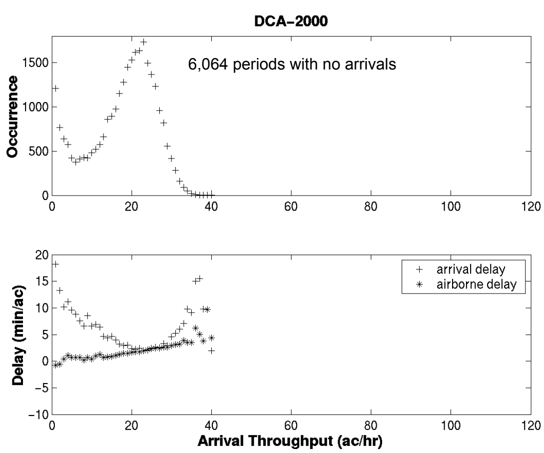

Figure 7. Spectra for Washington (DCA).

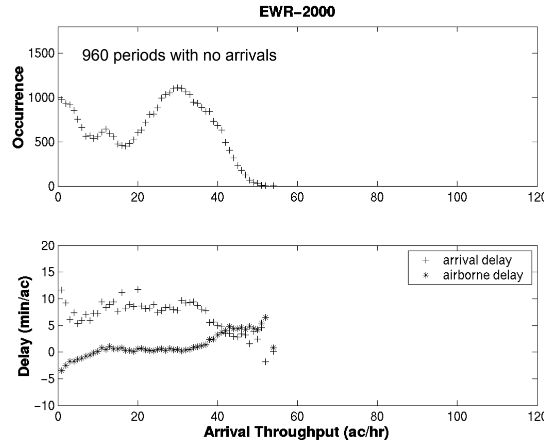

Figure 8. Spectra for Newark (EWR).

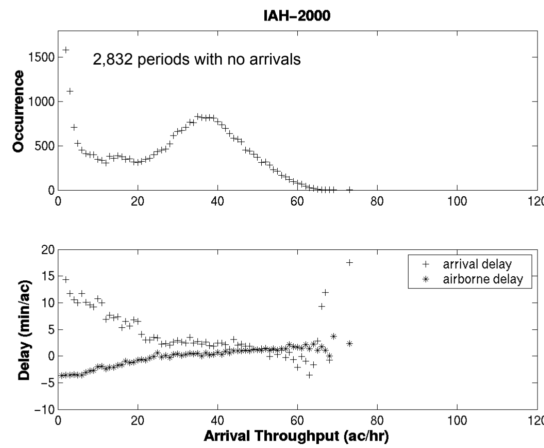

Figure 9. Spectra for Houston (IAH).

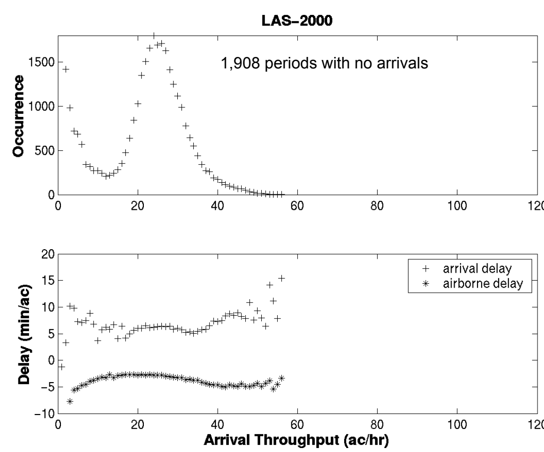

Figure 10. Spectra for Las Vegas (LAS).

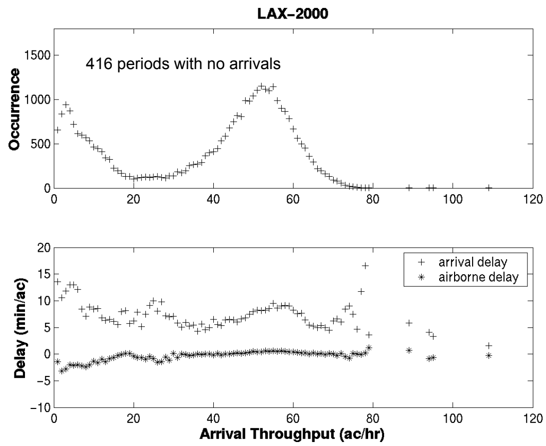

Figure 11. Spectra for Los Angeles (LAX).

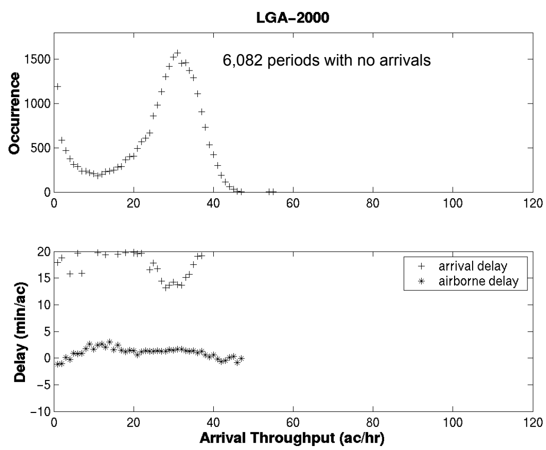

Figure 12. Spectra for LaGuardia (LGA).

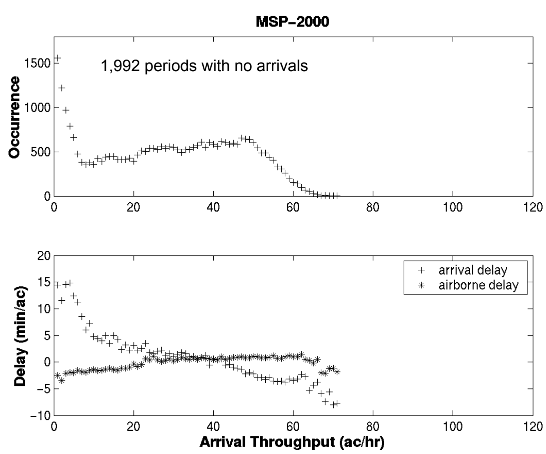

Figure 13. Spectra for Minneapolis/St. Paul (MSP).

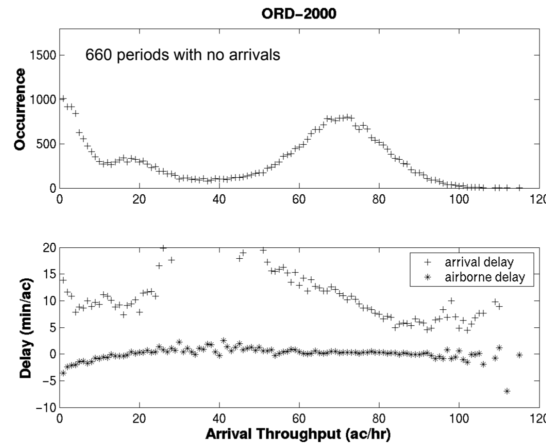

Figure 14. Spectra for Chicago O'Hare (ORD).

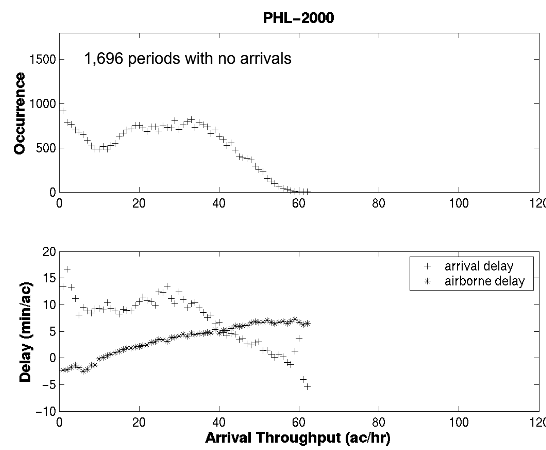

Figure 15. Spectra for Philadelphia (PHL).

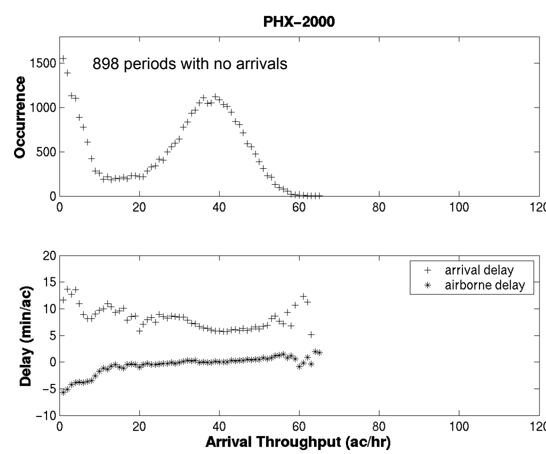

Figure 16. Spectra for Phoenix (PHX).

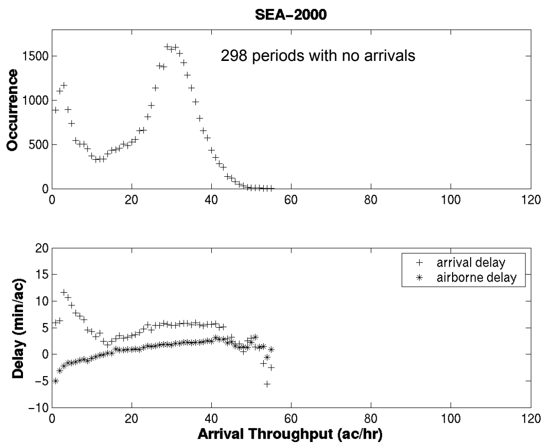

Figure 17. Spectra for Seattle (SEA).

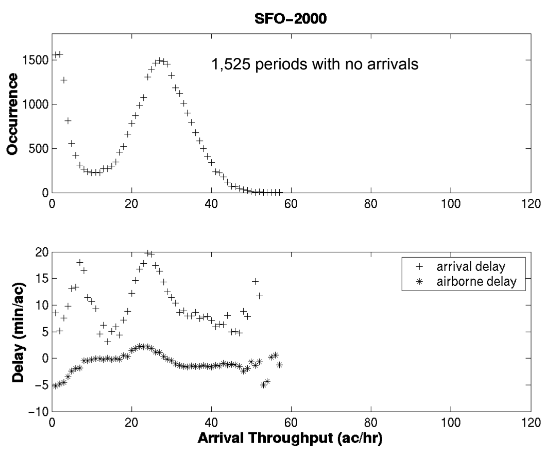

Figure 18. Spectra for San Francisco (SFO).

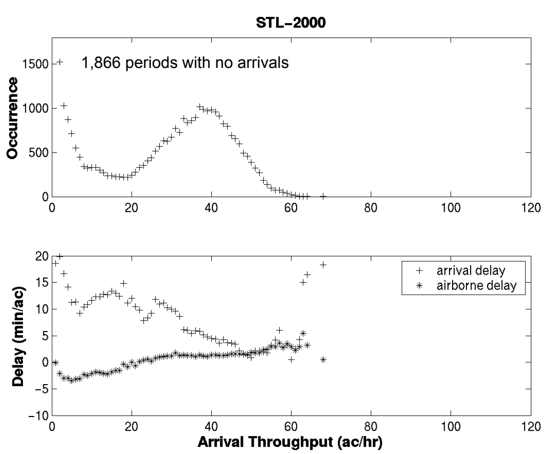

Figure 19. Spectra for Saint Louis (STL).

Most of the airports experienced a slightly greater positive slope in airborne delay from the lowest observed rate of 1 arrival per hour up to 10 or 20 arrivals per hour. BOS, DFW, EWR, PHX, and SFO experienced the most pronounced airborne delay increase at the low end of the spectrum. ATL and STL both reversed that trend with an initial *decrease* in airborne delay as the rate increased from 1 to 5 arrivals per hour.

The most significant airport-to-airport variations appeared in their arrival delay and occurrence spectra. Arrival delay at eight of the airports decreased with throughput. Other airports had oscillating arrival delay and some experienced U-shaped, peaked, and stepped delay patterns.

Las Vegas is the only one of these airports whose arrival delay generally *increased* over the entire throughput spectrum. Furthermore, Las Vegas is unique in that its average two-sided airborne delay was negative over the entire throughput spectrum. In a previous examination of airborne delay versus demand/capacity ratio, we found an inconsistency in the LAS demand, capacity, and delay data [6]. The FAA ASPM delay data and the FAA Benchmark Capacity data [7] indicate that LAS operates at 90% of capacity, but has the airborne delay of an airport operating at only 65% of capacity. The occurrence spectrum for LAS indicates that the peak-landing rate for the year was 56 arrivals per hour. This is consistent with the Benchmark Capacity data for LAS. LAS experiences most of its arrivals at rates between 20/hr and 30/hr, at about half its peak capacity. The fact that the average two-sided airborne delay is negative over the entire range of arrival rates at LAS indicates that aircraft operators systematically overestimate flight durations to LAS in their revised flight plans on takeoff.

Most of the airports had a distinct range of low occurrence counts between 10 and 30% of their peak throughput followed by a range of high counts (the distribution mode) between 40 and 70% of peak throughput. A sharply peaked occurrence mode indicates that most arrivals occur at or near the modal rate; that is, they are regularly separated in time. Regular arrival metering of this sort decreases the number of arrivals at higher and lower rates. On the other hand, single-carrier hubs, which experience significant variations in arrival rate over time, have relatively flat occurrence spectra. This behavior occurred at ATL, DFW, EWR, IAH, MSP, and PHL.

The effect of regular arrival metering on arrival delay is clearly observable. Airports with peaked occurrence spectra generally experience constant arrival delay over the entire spectrum. These airports seldom experience hubbing rushes, and as a result, queuing (airborne) delay generally does not rise as the arrival rate increases. Furthermore, at uniformly metered airports, arrival delay does not fall significantly at the highest extreme of the spectrum because instances of unusually high peak throughput can normally only occur when *delayed* flights arrive in bunches by chance. This is opposite to the situation at hubs, where the highest throughput normally occurs when aircraft arrive on schedule (that is, with low arrival delay) during hubbing rushes.

Arrival delay never remained below airborne delay over the full spectrum at any of the airports studied. At roughly half of the airports, the arrival delay was relatively uniform and remained larger than the airborne delay at all throughputs. The other airports experienced decreasing arrival delay that crossed and fell below the airborne delay at high throughputs. These crossovers occur because airlines adjust their schedules to account for average delay over the entire arrival spectrum.

A good quantitative predictor of these delay variations is the modal half-width of the occurrence spectrum, defined as the width of the modal throughput range at half of the mode count divided by the mode throughput. For example, at BOS the mode count was about 1400 occurrences and occurred at a rate of 34 arrivals/hour. The width of the distribution at the 700-occurrence level was about 14 arrivals/hour. The resulting modal half-width was thus 14/34 = 0.41. When the modal half-width is less than 0.5, the occurrence spectrum is highly peaked and the arrival delay distribution is relatively flat. Flat arrival delay spectra generally remain larger than airborne delay for all values of throughput. (SEA is an exception. Its CY 2000 modal half-width was 0.39, but its arrival delay dropped sharply below its airborne delay in the few periods for which the throughput was very high—greater than 55 arrivals per hour. This behavior repeated at SEA in CY 2001, and was likely caused by limited hubbing activity.)

Case Study: Effect of New Runway at PHX

Spectral analysis of airport delays provides a means for separating airport delays from en route delays. We used spectral analysis to examine the effects of opening the new runway at Phoenix International Airport on 5 October 2000. If one merely examines average monthly delay before and after a new runway opens, it can be difficult to detect the effect of the new runway because confounding factors can cause large variations in delay. Figure 20 compares the monthly averages of two-sided airborne delay for a seven-month period before the new runway opened and the same months in the year following the opening. The average delay increased at PHX after the new runway opened. The average delay from January to July 2000 was -0.47 minutes/arrival. The average delay from January to July 2001 was +0.36 minutes/arrival.

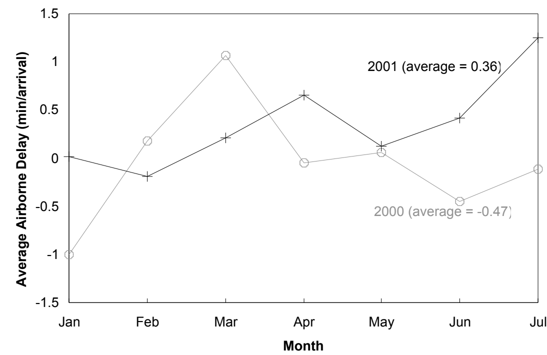

Figure 20. Delay at PHX over 7-month periods before and after adding the new runway.

When we compare the delay over 9-month periods directly before and after the opening of the new runway (we limit the periods to 9 months to avoid the effect of September 11, 2001), we also find an increase: from 0.41 minutes/arrival to 0.59 minutes/arrival. However, when we compare the delay *spectra* at PHX over those same 9-month periods before and after the opening of the new runway, we get a more complete picture and arrive at a different conclusion. Figure 21 compares the occurrence spectra and Figure 22 compares the two-sided airborne delay spectra for the periods before and after the runway opened.

Figure 21 shows that the peak throughput increased with the addition of the runway. In the nine months before the new runway opened, there were 109 windows with arrival rates exceeding the VMC benchmark capacity of 55 arrivals/hr at PHX. In the nine months after the new runway opened, the total arrival count increased by less than 2%, but the number of windows with arrival rates exceeding 55 arrivals/hr more than doubled to 285.

The FAA Benchmark Capacity Report [7] predicted that the VMC capacity would increase from 55 to 73 arrivals/hr when the new runway opened. A peak arrival rate of 66 arrivals/hr (90% of 73 arrivals/hr) occurred in two hourly intervals in the nine months following the runway opening. The mode of the distribution remained unchanged at 39 arrivals/hr, while the mode half-width increased from 0.49 to 0.62 arrivals/hr, suggesting that arrival metering was applied less frequently after the new runway opened.

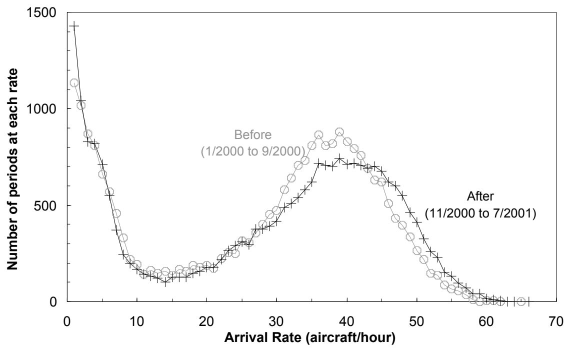

Figure 21. Occurrence spectra at PHX before and after new runway in October 2000.

Figure 22 shows that the opening of the new runway was associated with a decrease in airborne delay at high throughputs and an increase in airborne delay at throughputs below the mode (39 arrivals/hour). Queuing theory tells us that delay at arrival rates below about 50% of the airport capacity is less likely to be caused by terminal area queuing. However, terminal queuing is a significant component of airborne delay at high arrival rates.

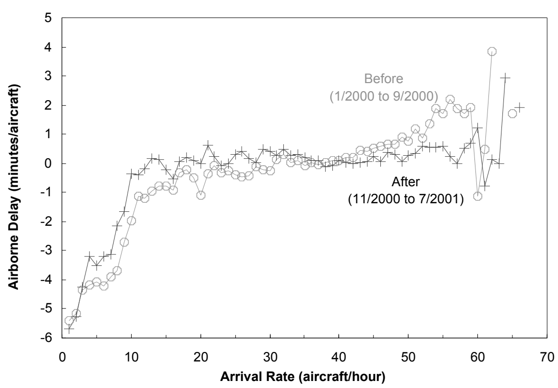

Figure 22. Airborne delay spectra at PHX before and after new runway in October 2000.

We also examined arrival delay at PHX for the same two periods. The only significant change in the arrival delay distribution occurred at the very lowest end of the throughput spectrum below 10 arrivals/hr. As with airborne delay, it is not likely that local runway capacity limitations directly caused arrival delay at rates below 10 arrivals/hr.

In addition to selecting the periods for this analysis to avoid the confounding effect of September 11, 2001, we also examined the effect of the general economic decline that began in 2000 before the new runway opened. Between August 2000 and August 2001, the Standard and Poors index dropped by about 20%. However, the demand at PHX in August 2001 was only about 2 percent less than the demand in August 2000, while the demand for all 32 Benchmark airports increased by about 0.2 percent, and overall US passenger enplanements increased by about 4%. Evidently, the stock market decline did not reduce air traffic demand between August 2000 and August 2001. To further assure that the observed improvements at PHX were not in some way related to general economic influences, we examined before and after spectral charts covering the same two 9-month periods for 15 other high traffic intensity airports. None of the other airports experienced a comparable combination of throughput and airborne delay improvements at high arrival rates after October 2000.

In the absence of any obvious confounding factors, the association of the new runway with increased throughput and decreased airborne delay at high throughputs suggests that the new runway improved performance at PHX.

Conclusions

Spectral analysis of airport delays provides a simple, but powerful filtering tool for separating airport delays from en route delays. This separation is possible because terminal queuing is the largest component of queuing delay at arrival rates approaching the airport capacity, whereas delay at low arrival rates is usually associated with en route effects.

Spectral analysis clarifies operational distinctions between airports by highlighting arrival counts and delay slopes at key points in the throughput spectrum. The resulting information would be difficult to extract from time-series analysis of aircraft arrivals. Such graphic distinctions facilitate airport classification. By sampling airports that belong to distinct spectral groupings, simulations and benefits studies may be able to generate more efficient and conclusive results.

Airport spectral analysis can also help clarify the effects of airport improvements by comparison of before and after spectral characteristics. Examination of delay spectra might also help aircraft operators improve scheduling accuracy.

This paper examined spectra for a number of important US airports for calendar year 2000. At all of these airports, delay relative to predicted flight time (airborne delay) was nearly invariant to arrival throughput at all but the lowest arrival rates. This indicates that airborne delay at intermediate throughputs occurred mostly in en route airspace and not at the arrival runways.

In contrast to airborne delay, delay relative to schedule (arrival delay) varied significantly with arrival throughput. At airports with well-metered flow, arrival delay generally increased at the high end of the spectrum. It is likely that the highest throughput occurred at these airports when delayed flights arrived in bunches by chance. On the other hand, at hub airports, arrival delay decreased at the high end of the spectrum, probably because instances of high peak throughput occurred more frequently when aircraft arrived on schedule during hubbing rushes.

References

[1] Welch, J.D., Andrews, J.W., and Robinson, J. E., "Assessing Delay Benefits of the Final Approach Spacing Tool (FAST)," AIAA Guidance, Navigation & Control Conference, Montreal, (Aug 2001).

[2] Allen, S. et al, "Analysis of Delay Causality at Newark International Airport," 4th USA/Europe Air Traffic Management R&D Seminar, Santa Fe, NM, (December 2001).

[3] Federal Aviation Administration Office of Policy and Plans, www.apo.data.faa.gov. (April 2001)

[4] Andrews, J. W. and Robinson, J.E., "Radar-Based Analysis of the Efficiency of Runway Use," AIAA Guidance, Navigation & Control Conference, Montreal, (Aug 2001).

[5] Welch, J and Lloyd, R., "Estimating Airport System Delay Performance," 4th USA/Europe Air Traffic Management R&D Seminar, Santa Fe, NM, (December 2001).

[6] Welch, J.D., "Assessing Capacity Needs of Congested Airports," AIAA Aircraft Technology, Integration, and Operations Forum, Los Angeles, (October 2002).

[7] Federal Aviation Administration Office of Policy and Plans, "Airport Capacity Benchmark Report 2001," www.faa.gov/events/benchmark (September 2001)

Key Words

Delay, capacity, airport, runway, spectrum.

Biographies

Jerry D. Welch

Dr. Welch is a Senior Staff Member in the Lincoln Laboratory Air Traffic Control Systems Group, and was formerly leader of the Air Traffic Automation Group where he lead efforts to develop oceanic and en route automation and radar-driven runway hazard alerts. He is currently supporting NASA in the analysis of computer aides for air traffic management. He was part of the team that developed the Mode S air traffic control radar beacon system for the FAA. He led the program at Lincoln Laboratory to develop the surveillance component of the Traffic Alert and Collision Avoidance System (TCAS). He helped the FAA initiate the Terminal Air Traffic Control Automation Program that led to the Center TRACON Automation System (CTAS) activity at NASA.

Shafique Ahmed

Dr. Ahmed is an Associate Staff Member in the Lincoln Laboratory Air Traffic Control Systems Group. He participated in the development of the Integrated Terminal Weather System for Air Traffic Control (ITWS) while working in the Weather Sensing Group at Lincoln Laboratory. He also worked in the Signature Studies and Analysis group at Lincoln Laboratory. He received his M.S. in Physics from Islamabad, Pakistan and Ph.D. in Physics from Warsaw, Poland. He also received a M.S. degree in Computer Science from Rensselaer Polytechnic Institute.