USING HISTORICAL FLIGHT DATA TO EVALUATE AIRBORNE DEMAND, DELAY AND TRAFFIC FLOW CONTROL

Michael Brennan and Terence R. Thompson, Metron Aviation, Herndon, VA, USA

S. Bradford and D. Liang, FAA/ASD-100, Washington, D.C., USA

Abstract

Efforts to fully understand the relationships among air traffic volume, system capacity and airborne delay has been a continuing field of research in the international Traffic Flow Management (TFM) community. These efforts differ in the ways in which airborne demand is modeled and quantified, and in how delay is measured. This paper explores the application of a neglected source of historical data that can be used to provide new insights into air traffic delay and Air Traffic Control (ATC) actions. The Enhanced Traffic Management System (ETMS) provides detailed information in real time about the path and en route delays taken by every flight in the US airspace, and several years of such data has been archived. Properly organized and applied, this information can yield unique and valuable quantification and visualization of the evolution of demand in the airspace; the occurrence, distribution and propagation of airborne delays; and the effects of traffic controller actions. Such insights range from system-wide pattern identification and analysis to flight-specific behavior and results. This paper presents the results of some of the analysis based on this data that has recently been conducted.

Introduction

As better methods of managing terminal area demand/capacity imbalances come on line, and as airports continue to expand and increase arrival rates, limited capacity in the en route airspace is becoming more and more the limiting factor in aviation operations. Efforts have been underway for years to understand the dynamics of the airspace, with the goal of improving capacity and efficiency, and much progress has been made, but the need and opportunity for further improvements remain.

Most previous work has been constrained by the difficulty in identifying where and when en route delays occur. Little information is available historically on what ATC controls initiatives have been applied, any record of the affected flights is rarer still, and data on the exact location, timing or extent of the delays incurred is non-existent. Typically delay analysis has been limited to identifying flights with longer than expected times en route (ETE), and trying to find insight from the extension of ETE alone (see Hoffman and Voss [1]). But there exists a wealth of data from which precise information on the location and timing of airborne delays can be obtained, and this source is largely untapped.

In this paper we present some of the methodologies being used at Metron Aviation, Inc, in conjunction with the FAA's Office of System Architecture and Investment Analysis, to exploit this information source. We also present some of the initial results that have been obtained. The goal throughout the effort is to improve our understanding of the relationships among air traffic volume, TFM and ATC procedures, and airborne delay. Such knowledge will in turn lead to better TFM policies and to better decision making by controllers.

System-wide behavior can be observed though the aggregation of data on all flights, but sometimes specific issues can be better investigated by limiting the data universe to an appropriate subset. One such grouping used in this paper is flights bound for Chicago O'Hare airport (ORD). En route delays are often caused by flights competing for the same resources, and flights bound for a given airport typically share a number of critical resources, including congested sectors, waypoints and jet routes. Focusing on a single airport's flights better illustrates the flight interactions that lead to delays, ORD is of special interest to the TFM community because it is often the busiest airport in the US and is a hub for two major carriers. ORD-bound flights contribute heavily to en route congestion in the New York to Midwest corridor, and ORD is subject to thunderstorms in the summer and heavy snowfalls in the winter. Therefore, some of the analysis presented here focuses on this portion of the airspace demand.

Historical Flight Data

TFM and strategic control of traffic flow in the US is conducted by the FAA based on the data provided through the Enhanced Traffic Management System. The ETMS system, developed and operated by the Volpe National Transportation Systems Center, collects and integrates information on pending and active flights and

other conditions relevant to TFM and distributes this information, with a number of value-added enhancements, to the aviation community in real time.

The flight-specific data collected by ETMS comes from airline schedules, operator intent information, controller and other FAA messages, and radar/transponder reports. This information is processed and redistributed to FAA facilities, airlines, and other users in a variety of formats and presented through a number of software tools.

This data is very valuable for real-time operations, but also is a rich source of information for post-operations and historical analysis. The ETMS data is correlated, integrated and archived at MAI in Herndon, VA, and there used for analysis and evaluation. This archive is the data source for the analysis in this paper.

The component of the ETMS data used in this analysis is the flight track output string. This data string provides, each minute, a status report on every flight currently active in the NAS. The primary source for this flight data is the en route radar/transponder reports, or TZ messages. A TZ message going into ETMS includes data on the aircraft ID and its current position, altitude and speed. However, ETMS has available to it additional information on the flight that is incorporated into the flight track report. The most important addition to the flight track report for the purpose of this analysis is an updated estimate of the arrival time for the flight at the destination airport, or ETA. How this field value is computed and how it is used are discussed in the section below on En Route Delays.

Airspace Demand

To understand and learn to manage congestion in the airspace requires tools to quantify and visualize the spatial and temporal distribution of en route demand. The ETMS flight track data provides the underlying data source that describes the en route volume. The total size of this data is large: 40.000 flights per day, one update for each per minute, 120 minutes average per flight produces almost 5,000,000 track reports per day. So some simple tools are needed to make the data accessible and useful.

One effective way to organize the track data for analysis and visualization is to integrate each flight track over a spatial and temporal grid constructed over the NAS. For these analyses, a grid over the continental US was defined, with position cells of ½ degree by ½ degree, temporal cells of 5 minutes duration, and altitude ranges of above and below 20,000 feet. The number of aircraft in each such cell averaged over the 5-minute lifetime of the cell can be computed from the track data, and this

then provides a detailed description of the en route volume.

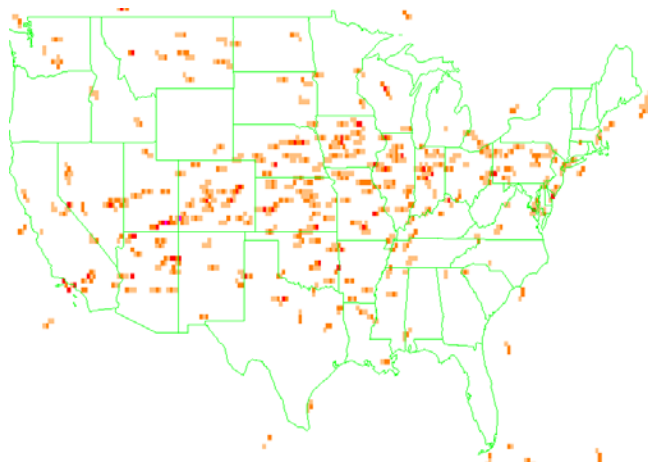Figures 1 through 4 illustrate this depiction of airborne demand

Figure 1: High Altitude Volume at 0800Z

The period around 0800Z is the quietest part of the aviation day in the US. Figure 1 shows the total volume above 20000 feet for the period from 0800Z to 0805Z on a typical day, September 27, 2002, plotted on a logarithmic scale. The main activity is from overseas flights and a few overnight cross-country flights.
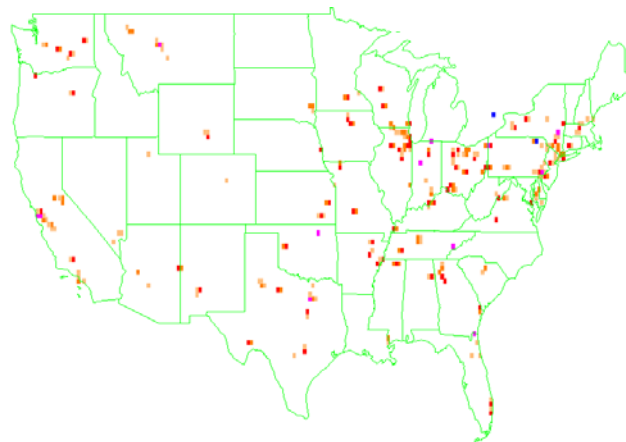

Figure 2: Low Altitude Volume at 0800Z

Figure 2 depicts the activity below 20000 feet at the same time. There is only light activity, primarily at or near the major airports.
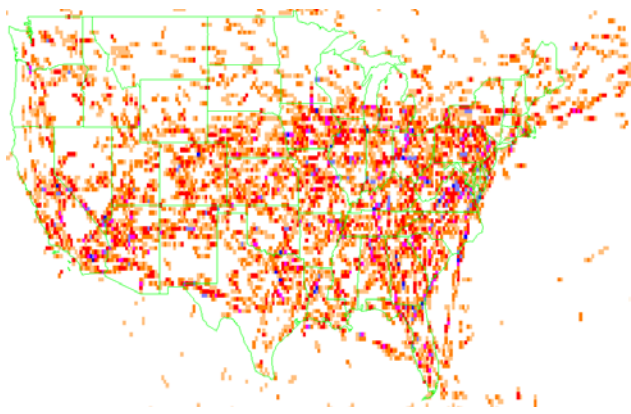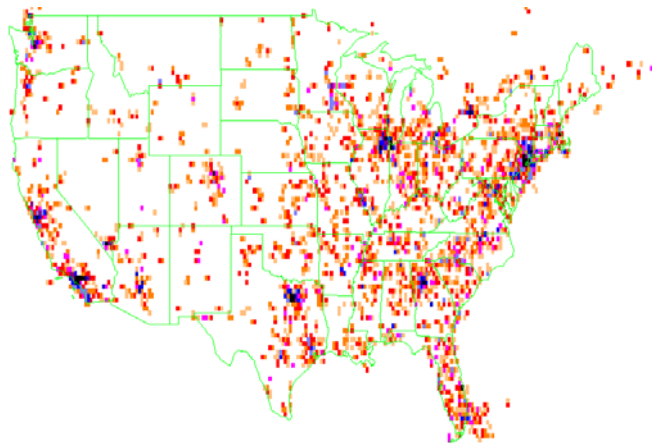
Figure 3: High Altitude Volume at 1800Z

Figure 4: Low Altitude Volume at 1800Z

Figure 3 and Figure 4 show the airborne activity 10 hours later. At high altitudes, volume is high over much of the NAS, highest in the triangle between the Northeast, Chicago and Atlanta. The heavily traveled air routes between key city pairs are clearly identified. At low altitudes there is very high volume around each major airport.
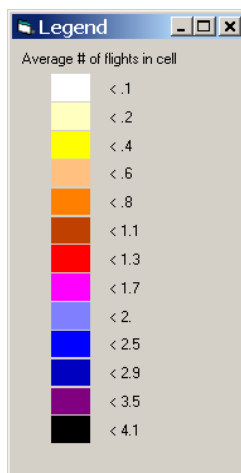

Figure 5: Color Scale for Volume Plots

Figure 5 shows the color scheme used in the volume plots. The densest cells have on average 4.1 flights in them over the 5 minute period.

En Route Delays

The goal of effective TFM and ATC, after ensuring safety, is to support the efficient flow of aircraft through the airspace. Ineffective TFM and ATC can lead to unnecessary en route delays for flights. To identify where and how these functions can be improved requires knowing where, when and, if possible, why en route delays are being taken. While pilots and controllers know in real time when flights are delayed, this information is not directly recorded. To perform any historical analysis of the effectiveness of TFM and ATC operations requires a way to identify, quantify and localize airborne delays from the available historical data.

Controllers impose delays on active flights to reduce demand on a limit resource somewhere up stream from the flight. The congestion point might be constrained arrival rates at an airport, a sector that is over capacity, or an area made impassable by severe weather. A controller has a number of tools that can be used to delay a flight: the flight's ground speed can be reduced; it can be told to 'vector,' following a zigzag pattern to reduce the effective speed; it can be put into a circular holding pattern; or it can be rerouted.

It is very difficult to infer when a flight took unexpected delay simply by examining the flight's track. A delay suggests that the flight is not making the progress towards its destination that was expected. New airborne delay could be inferred from comparing the expected progress of the flight to the actual progress, as measured in part by the evolution of the flight's position. But the expected progress depends on the planned route of flight, the aircraft's performance characteristics, and the local weather. It would be completely impractical to use the historical data to compute where each flight should be at each minute of its path, even if the flight plan, aircraft performance and local weather were available.

But what is impractical to do historically is done in real time by ETMS and recorded as part of the flight track data. ETMS incorporates a Trajectory Prediction model that computes the expected path of the flight based on its flight plan, performance characteristics, and current wind conditions. This projected path is then used in real-time to predict the time of specific future events, in particular when the flight will cross the boundaries of each airspace sector in its path. The final event for the

flight is its arrival at the destination airport, and the predicted time for this event is the ETA.

Whenever ETMS receives a new position report or amended flight plan for an airborne flight, it re-computes the times for each future event. If there is a change in the predictions the new event times are recorded. The historical track data thus provides a record of the evolution of each flight's ETA. So when a flight takes airborne delay its progress towards its destination will be less than expected, and this will be reflected in an increased ETA. The changes in the ETA predictions for a flight can then be used historically to infer where and when a flight is delayed en route.

This flight-specific delay data can be folded into the same grid structure used to track airborne volume by computing the spatial and temporal cell the flight was in when its ETA changed, and aggregating the delay over all flights in the cell. Figures 6 and 7 show the spatial distribution of total new airborne delays for the period from 1800Z to 1805Z
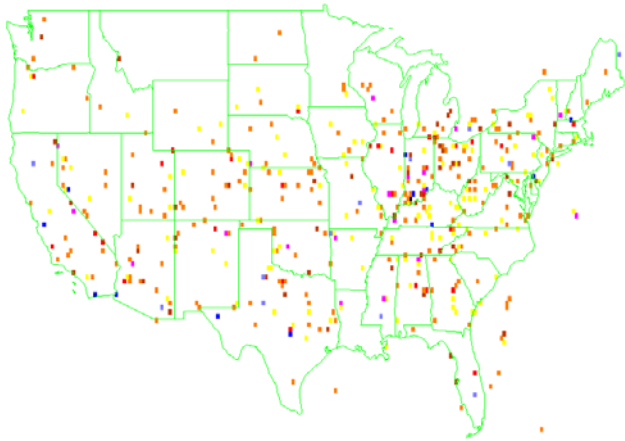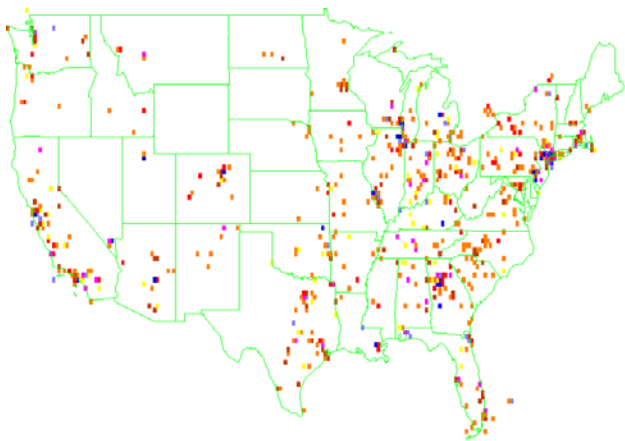

Figure 6: Airborne Delays at High Altitudes

Figure 7: Airborne Delays at Low Altitude

Figure 6 shows that the high altitude airborne delays are distributed fairly smoothly across the well traveled portions of the airspace. The distribution of delays is consistent with the distribution of traffic.

Figure 7, low altitude delay, shows intense concentrations around the busiest airports (EWR, LGA, ATL, ORD, DFW, LAX and SFO), suggesting the concentration of flights at these airports leads to long delays. But Figure 8 shows not the total delay but the average delay per flight. This illustrates that during this time period the high volume airports have the capacity to accommodate the demand.
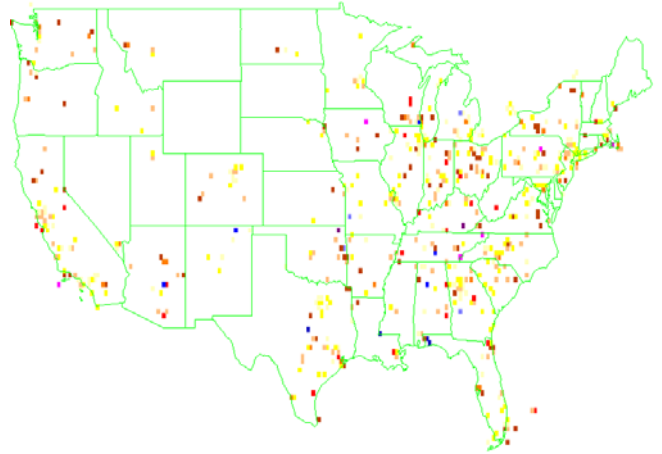

Figure 8: Average Airborne Delays at Low Altitudes

Figure 9 shows the color scales for the total and average delay plots. In the most extreme case flights are taking 15.8 minutes of delay when passing through the cells.
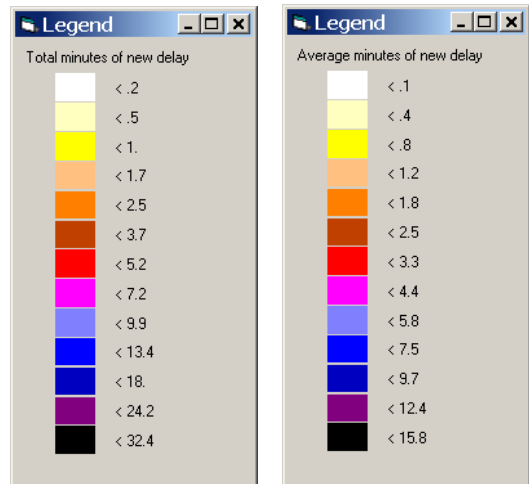

Figure 9: Color Scales for Delay Plots

The Relationship Between Volume and Delay

We do not expect to see simple relationships between the volume and delay distributions. For instance, there would be no natural correlation between the volume in a cell and the delay taken there – delaying a flight in a high volume cell would increase the volume in the cell, compounding the problem. Instead, delays are taken upstream from the congested area. The geographic plots of volume and delay by themselves are limited in assessing the volume/delay relationship because they miss a critical aspect of the dynamics of the system, the path taken by each flight through the airspace.

This suggests the need for a different way to visualize the volume and delay that recognizes the motion of flights through the airspace. Figure 10 is a plot of the path of a Memphis to ORD flight through the volume lattice. The horizontal axis tracks the spatial cells it moves through on its route. The horizontal mid-line of the plot shows the density in each cell at the time the flight passes through the cell. The vertical dimension shows the volume in each spatial cell before and after the flights passage. Cells below the center line represent earlier times (at five minute increments) and cells above reflect later times. The circles along the center line indicate times where the flight took airborne delay. The area of the circle is proportional to the amount of delay. The diagonal lines tracking down from each circle trace the volume in each cell on the flight's future path, not at the time the flight entered the cell, but rather at the time the flight took the delay. This visualization can be used to determine the degree to which controllers are basing their delay decisions on the upstream volume that will have materialized at the time the flight arrives at a point or on the volume ahead at the time the decision is made. This can help quantify the need for better predictive tools for controllers. In this example we see that although the flight was delayed en route, this control was not effective in that the flight still passed through the two highest volume cells in the window.
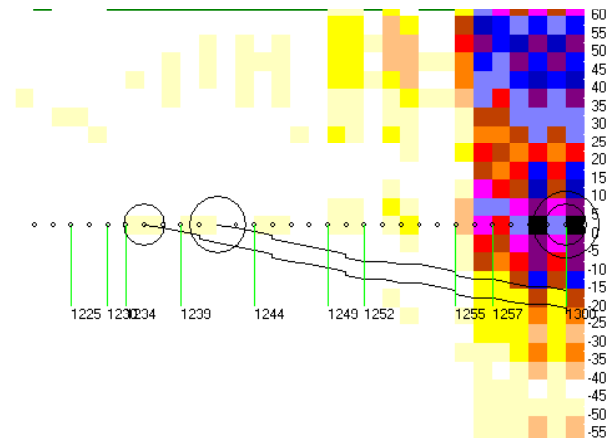

Figure 10: Delay and Volume History for a Single flight

An Example of Propagation of Delay in the Airspace

Excess demand at a terminal or in the airspace leads to airborne delay as controllers hold back flights to control demand. Holding flights in itself increases the demand in the airspace where they are held, and this in turn can lead to more holding and further delays upstream. Understanding the patterns and causes of how delays propagate is an important part of evaluating and improving TFM and ATM procedures. Tools based on the information in the historical ETMS data can provide valuable insight into the nature of delay propagation and suggest better policies and illustrate the need for better real-time decision aids and procedures.

Figure 11 shows the cell-based airborne delay map for flights bound for ORD. Each frame displays a snapshot of the delay incurred by flights over successive 5-minute intervals. At 1955Z a pocket of delay begins to develop over an area extending southwest from ORD. Five minutes later, the delay has intensified in the immediate area, and the extent of the delay has begun to spread further out from the airport. In each successive frame, the delay affects flights further away from the initial point of delay. By 2015Z, ATC delays have cascaded along a path from ORD to southern California across four or five states.
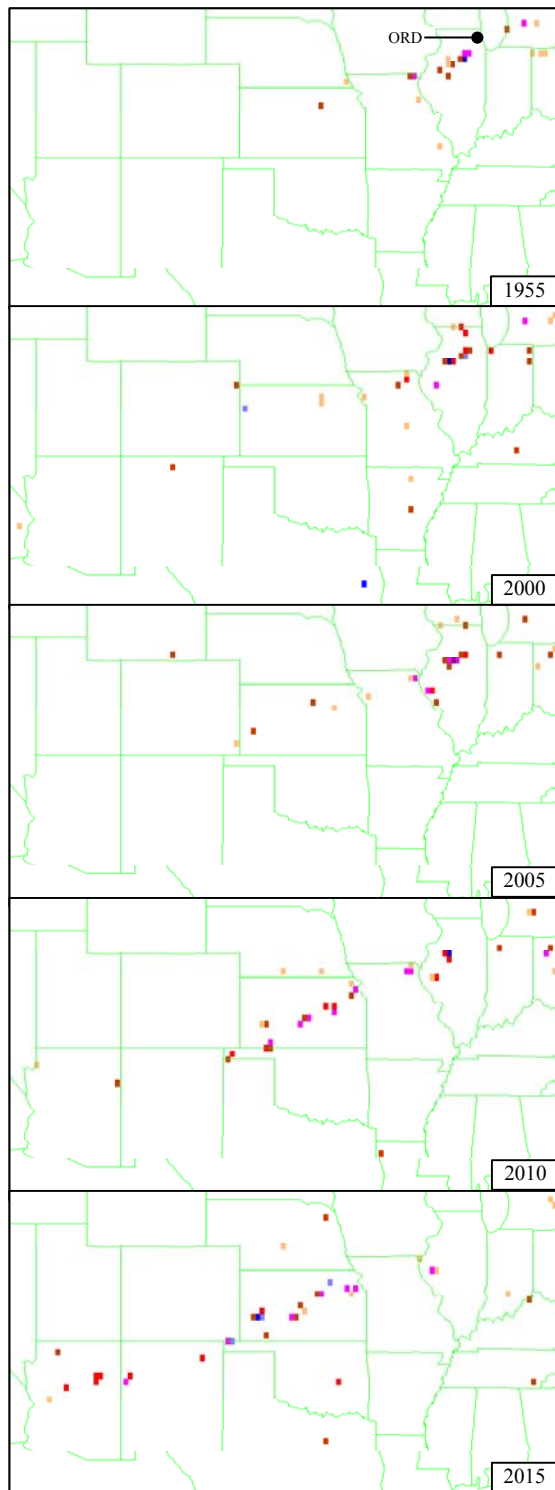
Figure 11: Propagation of Delay

In Figure 12 we see a detail of the delay map for ORD bound flights at 2015Z. Overlaid on the map are the tracks of the ORD-bound flights in the area for the 15

minute interval from 2010Z to 2025Z. The tracks for each 5 minute sub-interval are shown in red, blue and green respectively. The looping, S-turns and other maneuvers imposed on the aircraft by the controllers to reduce their effective speed are clearly evident. Black circles along each flight track show the position of the flight when ETMS inferred the slowed progress and reflected the delay by recomputing a later ETA. The area of each circle is proportional to the increase in the ETA.
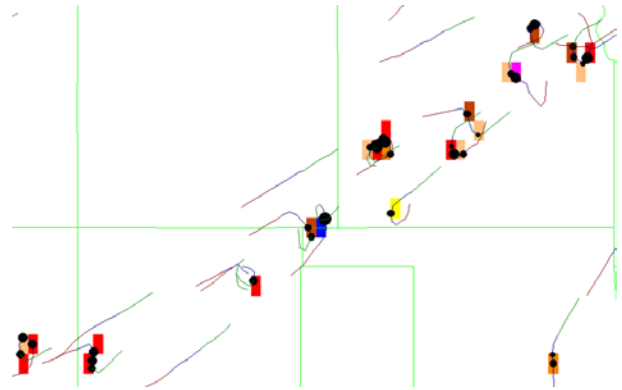

Figure 12: Flight paths for ORD-Bound Flights

A number of insights can be gained from these visualizations of the data. One is that the cascade of the delay is very rapid. It moves from the immediate terminal area out to over a thousand miles in 20 minutes. After the upstream delays have dissipated, flights are still being affected by the initial occurrence.

Another observation is that the imposed delay is largely restricted to flights over a particular jetway. Figure 13 shows the paths of all the ORD-bound flights for the two hours bracketing the time of interest. The root cause of this chain of delays was clearly excessive demand over ORD's southwest arrival fix, Plano. While there are several major streams converging on this fix, coming from Dallas, Houston, and New Orleans, virtually all of the delay is taken by flights on the jetway from Los Angeles and San Diego. This inequitable distribution of delay is no doubt a consequence of the limited tools available to controllers. The controllers' primarily tool to respond to such situations is to impose miles-in-trail restrictions on an airborne stream, and a single such initiative is easier to manage than multiple one. But the extreme maneuvers required to slow this single flow enough to sufficiently reduce upstream demand is a large burden to impose on the affected controllers and airspace users. Spreading the cost over a wider set of flights would be more equitable and over-all less painful. This suggests the need for better controller tools and a more even-handed policy.
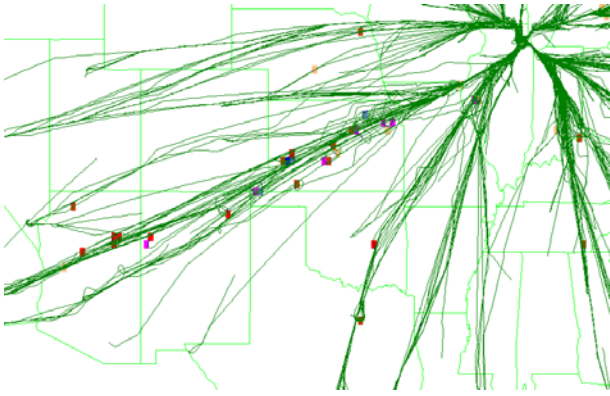
Figure 13: ORD-Bound Flight Streams

We see from this example that an overloaded arrival fix can create far-reaching and costly consequences for operators and controllers. Better methods for controlling demand at an arrival fix may have strong benefits. Procedures and tools to predict and control arrival fix demand through fix-specific Ground Delay Programs are currently being tested and evaluated at MAI (reference [3]).

Quantifying the Propagation of Delay

The example in the preceding section provides anecdotal evidence how of en route delay propagates, and provides the motivation to study this effect more systematically. In the example we see that the high delay at the original location dissipates fairly rapidly, while seeming to induce delay at later times to flights that are on track to pass through the high-delay cell. Some data manipulation and analysis will help show if this is unique incidence or a repeating pattern in the airspace.

To study this effect, we might consider a particular cell in our position/time grid where a high amount of delay was taken by flights. For convenience call the position cell X_0 and the 5-minute duration of the cell T_0.

Because we know the route of each flight we can also identify the flights that will be passing through X_0 at later times; for example, those that will be in the cell within 5 minutes of T_0, those that will be in the cell between 6 and 10 minutes of T_0, and so forth. That is, we can define flight-time based range rings around cell X_0. Call these range rings $\{R_i\}$. We can then compute for each range ring i $D_{i,0}$, the average delay taken by flights in that range ring at the time of the initial high delay T_0, as in the left side of Figure 14.
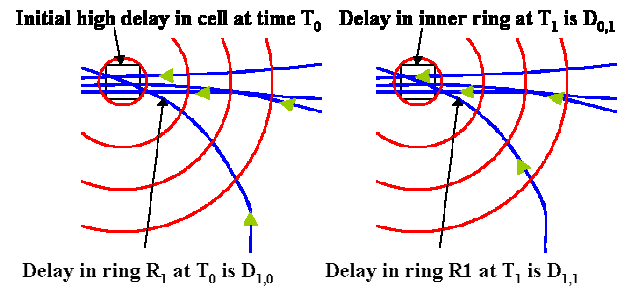

Figure 14: Range Rings and Delay Times

Additionally, we can determine, for each flight that will pass through X_0 at some time after T_0, the range ring that the flight will be in during each time interval T_i, and thus compute the average delay at distances out from the high delay cell at times after the high delay.

By combining such statistics from a number of cells where high delay was taken and averaging the values, we can construct a representation of the propagation of delay through time and space. Figure 15 depicts this relationship for one day's worth of flights in the NAS. The cell in the upper left corner of the figure represents the delay at X_0 and T_0 averaged over the high delay cells on that day (with the scaling adjusted to permit the showing of details in the other cells). The cell to its right shows the average delay in cell X_0 taken by flights in time interval T_1. Each subsequent cell on the right shows the average delay in cell X_0 5 minutes later. The top row in this figure thus gives us a measure of the temporal de-correlation rate of the delay, or how quickly the delay dissipates over time. We see that the delay drops off considerably over 10 minutes (and in this case later builds up a bit after 25 minutes). We can think of this as the propagation of delay through time.
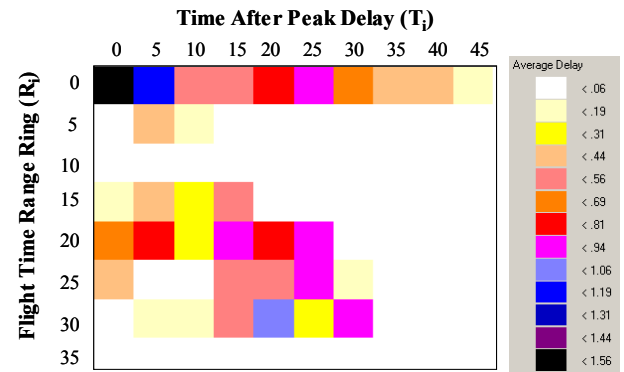

Figure 15: Spatial and Temporal Delay Propagation

The cell immediately below the upper left cell shows the average delay during T_0 taken by flights in each range ring around X_0; that is, it shows how flights headed for X_0 were delayed at T_0 as a function of their

distance from the cell. We can think of this as the propagation of delay through space.

The interior of the figure illustrates the propagation of airborne delay through both space and time. The shading in each cell indicates the average delay taken by flights as a function of their distance from the point of high delay (the cell's distance from the top row) and the time after the time of high delay (the distance from the left column). The highest delays (blue) in the interior are taken in cell (30,20), flights still 30 minutes flight time from the impacted cell took delays 20 minutes after the incident. The concentration of delay below the diagonal in the chart shows delays spreading rapidly over great distances.

These results are very new and their interpretation has just begun, but the important advance is the development of a methodology and the associated tool to investigate the general structure and specific instances of the temporal and spatial spread of airborne delay

The Effectiveness of Airborne Delays on Controlling Excess Volume

Under the premise that a principal motivation for imposing en route delays is to prevent over-saturation of any potion of the airspace, it is useful to verify that this effect is achieved. The ETMS data can be used to evaluate the effectiveness of delays taken by active flights in controlling excess demand.

The primary concern for controllers is the those few points in space and time when the volume is at its highest – their skills and tools allow them to safely and effectively deal with all but the rarest instances of high traffic volume. So the operational effectiveness of ATC actions can be measured by their success in reducing the instances and frequency of peak demand.

The distribution of airborne volume that results from controllers' actions is straightforward to compute: the empirical volume in each cell of the spatial and temporal grid defined earlier can be computed and profiled. But can it be determined whether the delays imposed on flights resulted in a better distribution of demand?

To do this, we need to compute what the volume distribution would have been had the delays not been applied. For this computation, we can consider the path of each flight, identify at what times and how much delay was taken, and recompute what the path would have been without the delay. That is, if at time t we see that a flight was delayed x minutes as measured by an x minute increase in its ETA, then we can model the undelayed track by presuming it would have been at each

point after t at a time x minutes earlier than it actually was. If further along its track, at time t', it takes an additional delay of x' minutes, then we can presume it would have been at each point after t' at a time $x+x'$ minutes earlier. We can then compute the evolution of the volume over the space/time grid based on the undelayed tracks, and use this as a model of the demand distribution in the absence of ATC controls.

Figure 16 shows the results of such an evaluation for one typical day of traffic. The graph is a histogram of the highest volume space/time cells over the course of the day. Traffic volume increases to the right. The vertical axis indicates the frequency of the occurrence of each volume level. The green bars, marked 'Delayed,' show the actual volume distribution, based on the recorded track data which includes the ATC delays. The red bars, marked 'Undelayed,' show what the distribution would have been in the absence of such delays. Clearly the frequency of the periods of highest demand has been significantly reduced by the ATC actions, reflecting well on the effectiveness of the current system.
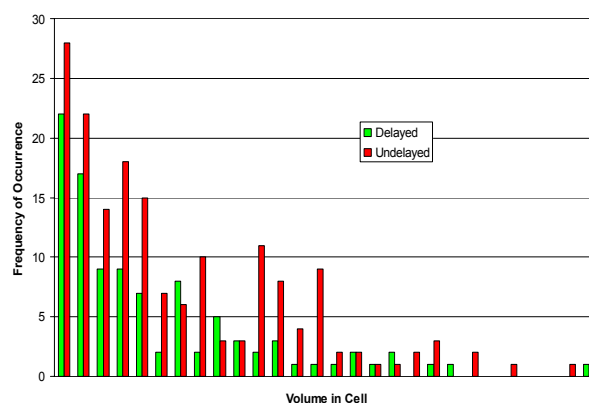

Figure 16: Histogram of Peak Volume with and without Delays

Figure 17 is a similar plot restricted to flights bound for ORD. The results are even more pronounced. Without the airborne delays, the maximum volume would have exceeded the highest realized volume frequently and severely.
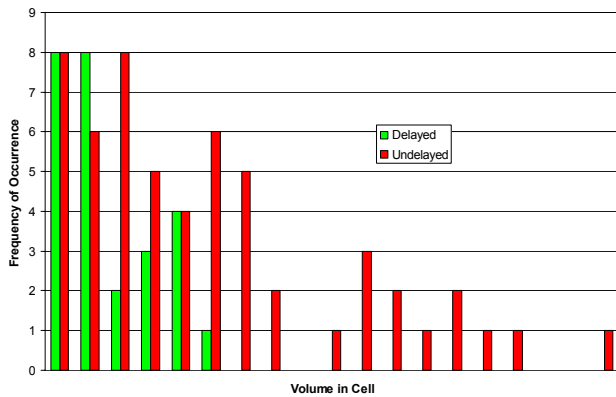
Figure 17: Histogram of Peak Volume with and without Delays for ORD-bound Flights

Conclusions

This effort to use the flight tracks and the en route delays inferred from the ETMS data to understand the relationships among volume, delay and air traffic operations has only been underway for a short time, but has already borne much fruit. Changes in the ETA estimates in the ETMS data has been shown to be an effective way of tracking airborne delays and in identifying the consequences of controller actions.

In these studies traffic volume was used as the measure of air traffic demand. It is well known that the workload is a more complex function of aircraft activity then simply flight counts. Future efforts will consider more realistic measures of air space complexity, drawing on the results of such works as Daniel Delahaye, Stéphane Puechmorel in reference [3].

Acknowledgements

References

[1] Joe Hoffman and William Voss, "Analytical Identification of Airport and Airspace Capacity Constraints", Third International Air Traffic Management R&D Seminar ATM-2000, Napoli, Italy June 2000

[1] ETMS System Design Document Version 6.0, section 25.{www1.faa.gov/tfmModernization/background-docs/misc-docs/ETMS-SDD/25-FDB.pdf}

[2] VNTSC-DTS56-TMS-002, Chapters 5 and 7.

[] Daniel Delahaye, Stéphane Puechmorel, "Air Traffic Complexity : towards Intrinsic Metrics", Third International Air Traffic Management R&D Seminar ATM-2000, Napoli, Italy June 2000

Key Words: traffic flow management, air traffic control, airborne delay, delay propagation, en route density, congestion

Biographies

Michael Brennan is the chief scientist for Metron Aviation, Inc (MAI), in Herndon, Virginia, and technical manager of the Technology and Infrastructure division there. He received a BS in Physics from the University of Pennsylvania and an MS in Computer Science and Operations Research from Villanova University. He is responsible for the conceptual development and validation of mid to near term TFM innovations at MAI. As chief scientist, he also oversees the coordination and technical quality of all of MAI's analysis and research efforts. Mr. Brennan has been involved in aviation for six years, and prior to that worked on analysis and decision aids for the Department of Defense.

Steve Bradford is the chief scientist for Architecture and NAS Development in the FAA's Office of System Architecture and Investment Analysis (ASD). In this role he has participated in the development of the RTCA NAS Operational Concept and the ICAO ATMCP Global Concept. His organization is also responsible for leading the effort to validate the future concepts, develop the FAA's ATC Information Architecture and leads several co-operative efforts via action plans with Eurocontrol. Prior to his current position, Mr. Bradford was the Manager of the NAS Concept Development Branch. Earlier, Mr. Bradford was lead on several simulation and analytic software development efforts, and conducted early analysis of Free Flight Concepts.

Diana Liang works for the Office of System Architecture and Investment Analysis for the Architecture and System Engineering Division. She is responsible for the development of the NAS Architecture Tool and Interface called CATS-I, directing analyses in support of NAS Concept Validation, and the development of Modeling Tools and Fast-Time Simulations to support that validation. This work includes several models she is developing jointly with NASA and cooperative efforts with Europe via Eurocontrol. Prior to working for ASD, Ms. Liang

worked in the Office of Energy and Environment for two years as the lead for the Emissions and Dispersion Modeling System (EDMS).