3rd USA/Europe Air Traffic Management R&D Seminar
Napoli, 13-16 June 2000

Closed Loop Forecasting of Air Traffic Demand and Delay

Peter F. Kostiuk
David Lee
Dou Long
Logistics Management Institute
McLean, Virginia USA 22102
Pkostiuk@lmi.org

Abstract

This report presents an integrated set of models that forecasts air carriers' future operations when delays due to limited terminal-area capacity are considered. This report models the industry as a whole, avoiding unnecessary details of competition among the carriers. To develop the schedule outputs, we first present a model to forecast the unconstrained flight schedules in the future, based on the assumption of rational behavior of the carriers. Then we develop a method to modify the unconstrained schedules, accounting for effects of congestion due to limited NAS capacities. Our underlying assumption is that carriers will modify their operations to keep mean delays within certain limits. We estimate values for those limits from changes in planned block times reflected in the OAG. Our method for modifying schedules takes many means of reducing the delays into consideration, albeit some of them indirectly. The direct actions include depeaking, operating in off hours, and reducing hub airports' operations. Indirect actions include the using secondary airports, using larger aircraft, and selecting new hub airports, which, we assume, have already been modeled in the FAA's TAF. Users of our suite of models can substitute an alternative forecast for the TAF.

1. Introduction.

This report describes a series of models that will forecast air carrier operations in the future as demand for air travel increases more rapidly than the ability of the National Airspace System (NAS) to accommodate flights in today's patterns. The work advances the discussion of NAS capacity beyond the level of recent dire forecasts that we are going to see massive flight delays in the United States and Europe within the next decade [1,2,3,8]. Those huge delays will not actually happen. They *would* occur if air carriers attempted to meet all future demand by simply increasing the number of flights while maintaining the same scheduling practices and other operating methods in use now. The airlines will certainly change their operating practices long before massive delays develop.

This report is concerned with predicting what those changes might be, and their likely economic implications. It is about the carriers' strategic response—not their tactical maneuvers—to operation disruptions caused by capacity-related flight delays.

In the *1998 Current Market Outlook,* Boeing recognizes the capacity limitations on air traffic growth and says that the carriers will find ways to increase traffic despite those limitations [4]. Specifically, Boeing suggests that airlines might take the following actions to avoid traffic congestion:

- Avoid congested hubs and gateways
- Use secondary airport metropolitan areas
- Move flights to off-peak time
- Broaden the range of departure times
- Shift short-haul flights to long-haul flights.
- Increase airplane size, especially for the short-haul flights.

While Boeing's list gives the airlines ways to grow the market under the capacity limitation, we also expect results from the severe congestion predicted. Those results are as follows:

- Increased flight block times or more padding in flight schedules
- Creation of new hub airports
- More slot-controlled airports.

While previous studies sometimes mention airline options for dealing with congestion, they have not incorporating those behavioral responses into the future demand forecasts. Consequently, analyses of the potential benefits of air traffic system improvements tend to be based on unrealistic demand scenarios and to generate implausible delay estimates. This report develops methods for quantifying the impacts of these alternative airline strategies in response to air traffic capacity constraints. The output of the approach is a schedule that explicitly accounts for system constraints and different airline responses to those constraints. These feasible schedules can then form the basis for a more credible assessment of future problems and policy evaluations. The models also predict how much of predicted growth in air travel is operationally feasible, given expected capacity limitations.

While the literature and applications are abundant on the topics of estimating the underlying variable distribution with constrained observations, we have not identified any similar work—either in academic publications or in airlines' practices—in developing air carrier operations models that incorporates delays. Some NAS software simulation packages, like NASPAC and AIRNET, have the capability to generate the future OAG schedule based on the current schedule and on the specified terminal growth rate, but they lack the capability to adjust the schedule due to excessive delay, which is the core of our present model. The NAS simulation packages may claim to have the capability to construct a new schedule for new hub cities, but they are essentially itinerary builders, and their users have to supply the hub cities' schedules and equipment types.

This paper describes the first successful effort to develop a methodology for closed-loop modeling of air traffic and aircraft flight delays. It therefore provides a more realistic framework for forecasting future demands on the ATM system and for evaluating the possible impacts of various airline strategic responses to system capacity constraints.

2. Modeling Constrained and Unconstrained Flight Schedules.

Our focal point is on the airline service schedules, which are described by the origin and destination airports, departure and arrival times, aircraft types, and operators. Airport operations are defined as the total number of departures and arrivals at an airport in a given time period, e.g., an hour. Given the flight schedules of all airlines, one can get the operations for any airport at any time by simply aggregating the appropriate flights. A completed flight from the origin gate to the destination gate requires many air traffic control (ATC) services, each of which may cause delays. Delays may be caused by shortages in taxi capacity on the ground; runway capacities, holding aircraft either on the ground or in the air; terminal radar approach control (TRACON) capacity in the air near the airport; or en route ATC capacities. Since the overwhelming majority of flight delays in the United States will be caused by the airports, either from runway or taxiway shortages, here we model only the impacts of limited airport capacities on air carrier operations [1,2,8].

We begin by exploring what the carriers will do in the future if they are rational. Suppose there are no excessive delays. Then they will conduct their operations as predicted by the unconstrained demand forecast model. The reason that they will have such schedules is that they are the most economical ones, absent excessive delay.

Now suppose that some flights at some airports at some times during the day have to be canceled due to excessive delays induced by excessive demand. The following assumptions seem to be reasonable for modeling an air carrier's operations due to flight delays:

1. As demand increases, the airlines will keep growing their schedules to meet it according to the unconstrained model, until the resulting operations produce excessive delays.
2. The delays chiefly important in airlines' schedule planning are those that occur in terminal areas due to insufficient runway and taxiway capacities in universally good weather conditions.
3. Any canceled flights must have the least economic value to the carriers. Since the variable cost to carry an additional passenger is so negligible in the flight operation, and most other costs must be distributed to the entire network, we are not going to consider cost explicitly in modifying the schedules. Instead, we focus on revenues.

Figure 1. A Model for Forecasting Constrained Flight Schedule in the Future

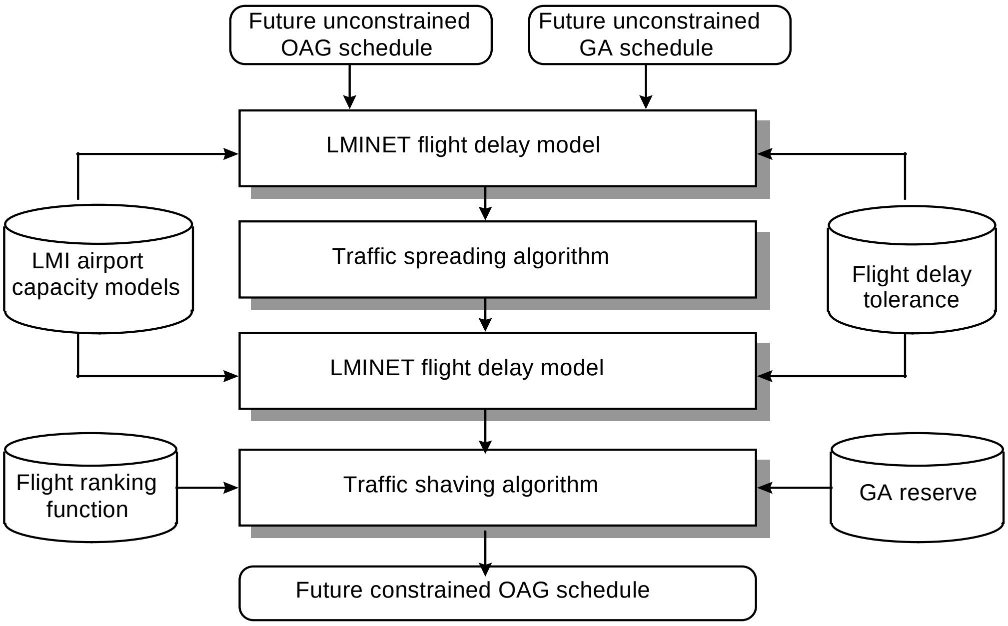

Figure 1 is a schematic of the model that we propose for forecasting air carriers' future flight schedules as they respond to increasing delays.

The following subsections detail the input data sources, individual components of the model, and their working and usage. A user is free to modify or use any other comparable data input for the model. For the sake of completeness, the model for forecasting the unconstrained flight schedule, which is part of the overall model, is included.

CURRENT FLIGHT SCHEDULE FOR AIR CARRIERS

This is a data input; we use the OAG schedule.

CURRENT FLIGHT SCHEDULE FOR GENERAL AVIATION

This is a data input; we use the schedule extracted from ETMS based on April 8, 1996.

TERMINAL AREA FORECAST

This is a data input that provides airport-specific traffic (operations) growth rates. We use the growth rates given in the FAA's TAF.

UNCONSTRAINED AIR TRAFFIC DEMAND FORECAST MODEL

For the air carriers, we expand the current OAG schedule using the TAF's airport-specific growth rates and the Fratar algorithm. For GA, we expand the current representative GA schedule, determined from ETMS data, using the TAF's airport-specific growth rates for GA and the simple algorithm of multiplying the departures by the GA growth rate of the departure airport. The reason that we do not use any sophisticated algorithm here is that GA flights do not follow any particular schedule, and they tend to leave for a limited number of airports from a given departure airport, thus making the Fratar algorithm unstable.

AIRPORT CAPACITY MODELS

This is a data input. For each of 64 LMINET airports, we need two sets of models: the runway capacity model and the taxiway capacity model. The runway capacity model is characterized by a piecewise linear Pareto frontier characterized by four parameters. The taxiway capacity model is based on the $M/M/1$ delay model and by matching the mean delays at each airport to those in the FAA's PMAC. (References [1] and [2] give details of these models.)

AIRPORT DELAY TOLERANCE

This is a data input. This parameter controls the maximum allowable demand. The user has the flexibility to specify the tolerance individually for each airport, and by departure or arrival. The tolerance can be imposed either as a maximum average delay or as a minimum probability that the delay will not exceed 7.5 minutes. (Air carriers and DOT usually use delays of more than 15 minutes to calculate the delay statistics, which splits to 7.5 minutes for both departure and arrival). The relationship between the maximum tolerable average delay minutes and the minimum tolerable on-time probability depends on the initial condition of the airport (whether the airport is busy or not during the last time period) and on the capacities of the runway and taxiways. Numerical calculations for the 64 LMINET airports show that the two tolerances share a fairly stable relationship, with an error of 1 to 2 percent of the on-time probability, which is shown in Table 1.

Table 1. Maximum Tolerable Average Delay Minutes Versus Minimum Tolerable On-Time Probability

Maximum Tolerable Average Delay Minute	5.0	4.0	3.0	2.0
Minimum Tolerable On-Time Probability	78%	85%	92%	98%

The overall on-time percentage for 1996 as reported by ASQP is 78.9. This statistic includes delays caused by many sources. What is desired here is the tolerance just for the delay caused by the insufficient airport capacities, which certainly must be higher than the statistics from ASQP. Based on Table 1, the maximum tolerable delay, system-wide, must be less than 5 minutes. In previous work, we derived the average padding in the schedule, based on the flights among the 29 FAA large airports, to be 3.83 and 2.71 minutes for departure and arrival, respectively. Since these paddings are estimated for all the flights, whether the airport is busy or not, the system-wide tolerance must be higher, because that tolerance must be based on the padding when the airports are busiest. On the other hand, the 29 FAA large hub airports, being the busiest in the United States, experience most of their operations while operating close to their capacities. The tolerances inferred from these airports should not be far from the right one for the overall network.

Accordingly, throughout this report when we present the illustration of this model, we use 3.83 and 2.71 minutes as the maximum tolerable average delays for departures and for arrivals, respectively. (Very likely, the levels of tolerance are different at different airports. Lacking data on which to base good estimates of airport-specific tolerances, however, we used common values for all airports.)

AIRPORT DELAY MODEL

This model computes the delay distributions for both departures and arrivals, based on demands and capacities. Appendix A details the development of our model, which is based on accurate numerical solutions of the exact equations for two-member tandem queues representing arrival and departure processes.

For our problem of finding the maximum tolerable demand such that the maximum average delay is T, our algorithm works as follows. Let x be the unconstrained demand; x_1, x_2, x_3 be successive approximations of the maximum tolerable demand to meet the target; and y_1, y_2, y_3 be the average delays associated with x_1, x_2, x_3, respectively.

Step 0: *Let* $x_1 = min(x, 0.9 \times min(runway\ capacity, taxiway\ capacity))$.

If $y_1 \le T$ and $x_1 = x$, stop.

Step 1: If $y_1 > T$, then let $x_2 = 0.95 \times x_1$.

If $y_1 < T$ and $x_1 < x$, then let $x_2 = min(x, 1.05 \times x_1)$.

Step 2: If $y_2 = T$ then stop; else apply the secant method to generate x_3,

$$x_3 = x_2 - \frac{x_2 - x_1}{y_2 - y_1}\left(y_2 - T\right), \qquad \text{[Eq. 1]}$$

If $x_3 > x$, then let $x_3 = x$. If $x_3 < 0$ then let $x_3 = 0$.

$x_1 \leftarrow x_2, \qquad y_1 \leftarrow y_2;$

$x_2 \leftarrow x_3, \qquad y_2 \leftarrow y_3.$

The solution of the secant method is the maximum allowable demand. It turns out that the secant method converges fairly quickly for our problem. We have at most needed three iterations in step 2 with 0.1 convergence bound, which corresponds to at most an uncertainty of 0.5 operation/hour in demand.

TRAFFIC SPREADING MODEL

When demand exceeds maximum tolerable demand, we adjust the schedule to "spread" excess demand into neighboring epochs. This procedure is based upon the result proved in Appendix C, that balancing arrival and departure operations across all epochs in a day will maximize the total number of operations.

If the demand in 1 hour exceeds the maximum tolerable demand, the excess demand—i.e., the difference between demand and maximum tolerable demand—is accommodated by spreading it evenly into the neighboring hours (1 hour before and 1 hour after), so that in each epoch all the flights have the same proportion of retained flights from the original demand and the same proportion of spread flights from the neighboring hours. Due to the nature of our problem, the number of flights in an epoch is not necessarily an integer.

FLIGHT RANKING FUNCTION

This is a data input that will rank the desirability of flights. It will be used by the traffic shaving module (described later) to cut flights when necessary to meet the maximum tolerable delay. Ideally, the desirability of a flight should be directly related to its revenue contribution, or fares and load factor, computed based on the historical data, but those data are not available to us. As stated before, they are the most closely guarded of all the air carriers' data. Instead, we use a surrogate, defined as the number of connection possibilities of a flight, which is the sum of the total arrivals to the departure airport 1 hour before the flight departure and the total departures to the arrival airport 1 hour after the arrival. The implication of this policy is that flights to or from small airports will be cut first. Flight connection is widely used by the market planning department of the carriers in forecasting passenger demand.

AIRPORT GA RESERVE

This is a data input per airport and per epoch. It is the number of reserved GA slots, which the GA demand cannot exceed. The unused GA slots cannot be used by commercial flights even if the GA demand is below the assigned number of slots. This is currently airport operating policy under slot control. For the model runs in the report, we have set the GA reserves to 0 throughout.

TRAFFIC SHAVING MODEL

Some flights have to be cut if the combined commercial and GA demands exceed the airport capacity. The flight with least desirability and the flight that is mutually least desirable from both departure and arrival airports will be cut first. While the least desirable flight is well-defined by the ranking function, finding it involves a search throughout the entire network on both the departure and arrival airports' lists.

If the departure demand exceeds the capacity at one airport and we cut all the excessive departure flights in order to meet the airport capacity, then it is quite possible that some additional departures from this airport will also have to be cut due to the cutting of arrival flights at some airports, which will cause under-utilization of airport capacity even with excessive demand. Another problem with cutting all the excessive demand from airport to airport is that the flight-ranking function will be significantly modified; the sequence in which the airports' excessive demand is cut makes a big difference.

We solved this technical problem by incrementally shaving off excessive demand throughout the entire network. First, we start off with a padding added to the airport capacities such that there are no excessive flights. Then the padding is reduced by an increment, the lowest-ranking excessive flights are cut off, and the fight-ranking function is updated. This process continues until the padding is 0.

This algorithm reduces the aforementioned problem to the minimum that the demand may be over-cut by a delta at the worst, and the flight-ranking function is current.

4. Results and an Application to Technology Infusion.

Clearly, by the examples shown in the previous sections, either some traffic growth has to be curtailed because of insufficient NAS capacity, or the capacity must be increased. Since it becomes harder and harder to construct new airports or augment existing airports by paving new runways or taxiways—due to the available land, financing, environmental regulations, etc.—the opportunities to improve capacity by physical improvements to the airports' capacities is quite limited. (We do include "certain-to-happen" projects from the airport improvement plan published by the FAA in our airport capacity models for 2007.)

Another way to improve the NAS capacity is by improving airports' efficiencies. To meet the national requirement for sustained air traffic growth, the NASA Aeronautics Enterprise has recently asked LMI to undertake a study to evaluate the contribution of the existing programs [2] Two NASA programs, the Terminal Area Productivity (TAP) and the Advanced Air Transport Technologies (AATT), comprise tools that will essentially either reduce the uncertainties of the aircraft position or enhance the airport utilization by better sequencing of flights and balancing runway usage. Both types of technologies, in addition to some other benefits in terms of safety or flight operating cost, will enlarge airport capacities.

To assess the tools' effectiveness, we mapped the operating characteristics of each set of tools into the parameters of the LMINET airport capacity models [1,2] This yields another set of runway capacity frontiers and another set of taxiway service rates. Figure 4 shows an example of a changed airport capacity Pareto frontier. The capacity Pareto frontier with the technology dominates the one without.

Nationwide for the 64 LMINET airports the daily operations in 2007 are 61,668 with ATM technologies versus 60,120 daily operations without ATM technologies, achieving 92.7 percent versus 81.3 percent of the potential traffic growth from 1996–2007. Table 2 summarizes the total daily operations for each of the 64 LMINET airports and for the overall 64 LMINET airports. We include the model results for unconstrained forecasts, i.e., with no capacity constraints allowed, and the constrained forecast with and without the technology enhancements. In our complete report, we also estimate RPMs, enplanements, airline operating costs, and fare yields.

Figure 2. Airport Runway Capacity Comparisons

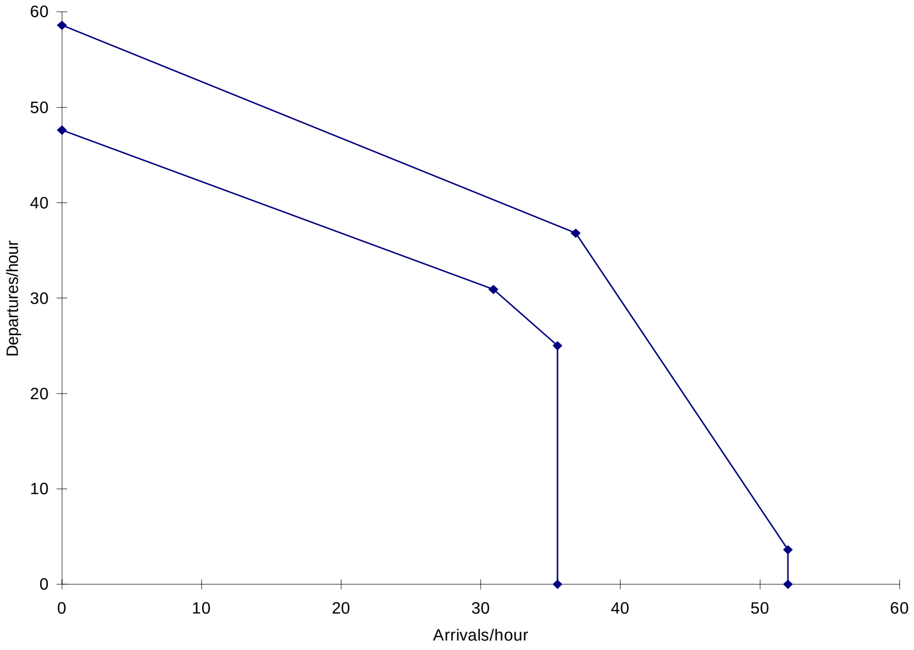

Table 2. Daily Carriers Operations by Airport

Airport	1996	Unconstrained 2007	Constrained 2007	Growth attained (%)	Constrained 2007 with tech	Growth attained (%)
BOS	1,262	1,374	1,371	97.4	1,374	99.6
BDL	347	431	422	89.9	428	96.8
HPN	165	209	204	88.8	208	96.9
ISP	119	150	149	97.4	149	99.4
TEB	0	0	0	100.0	0	100.0

Airport	1996	Unconstrained 2007	Constrained 2007	Growth attained (%)	Constrained 2007 with tech	Growth attained (%)
LGA	960	1,065	963	2.5	1,056	90.9
JFK	907	1,005	987	81.9	1,005	100.0
EWR	1,306	1,605	1,603	99.6	1,604	99.7
PHL	1,036	1,314	1,311	98.9	1,314	100.0
BWI	702	913	883	85.7	911	99.0
DCA	758	800	774	38.8	790	76.7
IAD	777	963	961	99.0	960	98.5
GSO	275	357	353	94.8	357	100.0
RDU	445	552	547	96.0	551	99.0
CLT	1,184	1,467	1,445	92.3	1,468	100.1
ATL	2,093	2,464	2,350	69.4	2,463	99.9
MCO	890	1,341	1,322	95.7	1,340	99.7
PBI	322	375	364	79.9	372	95.4
FLL	512	696	686	94.7	689	96.2
MIA	1,348	1,801	1,617	59.5	1,762	91.4
TPA	696	874	873	99.8	874	100.0
MSY	402	481	477	94.3	479	96.6
MEM	916	1,172	1,087	67.1	1,168	98.4
BNA	492	588	586	97.4	587	98.6
SDF	445	564	544	83.5	558	95.2
CVG	1,135	1,711	1,600	80.7	1,660	91.2
DAY	312	361	348	73.3	359	96.4
CMH	446	575	572	97.2	573	98.1
IND	567	788	786	99.0	787	99.4
CLE	798	1,033	1,030	98.7	1,033	100.0
DTW	1,351	1,815	1,742	84.3	1,814	99.9
PIT	1,220	1,470	1,468	99.4	1,467	98.8
SYR	285	353	347	90.7	353	100.0
MKE	460	584	563	83.3	582	98.8
ORD	2,367	2,793	2,698	77.7	2,793	100.0
MDW	479	608	599	93.3	607	99.2
STL	1,414	1,753	1,618	60.1	1,744	97.5
IAH	1,053	1,495	1,415	82.0	1,488	98.5

HOU	397	487	466	77.1	449	57.6
AUS	303	406	392	86.1	403	97.1
SAT	363	487	479	93.2	479	93.3
DAL	377	523	523	100.0	515	95.1
DFW	2,443	3,372	2,575	14.2	2,778	36.0
MSP	1,257	1,611	1,524	75.3	1,585	92.5
MCI	548	669	667	98.4	669	99.9

Table 2. Daily Carriers Operations by Airport (Continued)

Airport	1996	Unconstrained 2007	Constrained 2007	Growth attained (%)	Constrained 2007 with tech	Growth attained (%)
DEN	1,148	1,391	1,390	99.7	1,391	100.0
ABQ	366	490	487	97.4	486	96.5
ELP	206	233	225	73.0	230	90.3
PHX	1,312	1,779	1,632	68.6	1,736	90.8
SLC	830	1,138	1,129	97.1	1,136	99.4
LAS	1,024	1,510	1,506	99.2	1,510	99.9
SAN	611	798	750	74.3	786	93.4
SNA	303	418	402	86.4	414	96.9
LGB	38	54	40	13.6	44	35.8
LAX	2,042	2,517	2,457	87.3	2,487	93.8
BUR	289	391	386	95.3	384	92.9
ONT	362	442	434	89.4	437	93.3
RNO	295	414	410	96.7	411	98.0
SMF	354	489	480	93.3	484	96.4
OAK	698	887	877	94.7	884	98.7
SFO	1,111	1,443	1,422	93.8	1,442	99.6
SJC	363	506	499	94.8	501	96.2
PDX	731	987	986	99.3	987	99.9
SEA	1,062	1,319	1,318	99.6	1,318	99.6
Total	49,073	62,656	60,120	81.3	61,668	92.7

5. Conclusions.

This report presents an integrated set of models that forecasts air carriers' future operations when delays due to limited terminal-area capacity are considered. The suite has two outputs. The more detailed output consists of flight schedules, which convey much useful information about the air carriers' operations, including origins, destinations, and planned block times. The schedules are made by models that restrict traffic growth to levels that the NAS can accommodate with not more than user-specified values of mean arrival delay and departure delay per flight.

The other output is forecasts of commercial RPMs, enplanements, and total operations for the entire U.S. passenger air carrier industry. These results are made by linking econometric models to a NAS model to determine operations levels that will generate user-specified profit levels under delay-induced reductions in productivity.

This report models the industry as a whole, avoiding unnecessary details of competition among the carriers. To develop the schedule outputs, we first present a model to forecast the unconstrained flight schedules in the future, based on the assumption of rational behavior of the carriers.

Then we develop a method to modify the unconstrained schedules, accounting for effects of congestion due to limited NAS capacities. Our underlying assumption is that carriers will modify their operations to keep mean delays within certain limits. We estimate values for those limits from changes in planned block times reflected in the OAG.

Our method for modifying schedules takes many means of reducing the delays into consideration, albeit some of them indirectly. The direct actions include depeaking, operating in off hours, and reducing hub airports' operations. Indirect actions include the using secondary airports, using larger aircraft, and selecting new hub airports, which, we assume, have already been modeled in the FAA's TAF. Users of our suite of models can substitute an alternative forecast for the TAF.

Users can modify other features of our schedule-generating suite. In addition to the TAF forecasts, the parameters users are most likely to want to change are the airport delay tolerances and the airport capacity models. The users may also want to integrate their own individual models into the suite. However, the overall suite's present configuration is sound, and it will give defensible results efficiently.

GA and commercial flights are assumed to have equal relative flight reductions in the economic models, while GA is assumed to be cut first in the flight-specific models. Also, all the flights are assumed to be reduced equally in the economic model, while only flights in the congested airports are cut in the schedule-producing, flight-specific models. This means that more flights are cut in the economic models, since the delays are caused by intense operations in the hub airports.

No matter which models we use, either flight-specific or economic, the methods we have developed have profound implications for the evaluation of ATM technologies and procedures. First, it is likely that we are not going to see the dramatic increases of air traffic delays as predicted by other planning models, because of the airlines' adaptations to the constraining capacities. We see early signs of the validity of our approach and results because the seriously increased delays forecast by some researchers a decade or so ago have not materialized.

The second implication is based on the first, that we need new methods to evaluate the benefits of new ATM technologies and procedures. The model suites presented here—based on carriers' limited tolerance for delays and on their economic incentives to change operations in the face of productivity lost due to delays—are an initial response to this.

The models, accounting rationally for airlines' reactions to delay, are likely to be reasonably robust. This is not necessarily true for this report's specific detailed forecasts. Many uncertainties, some quite small, may change air carriers' operations in the future. We base our conclusions on the assumptions of rational behavior of the carriers and a relatively stable operations environment. Changes in the legal framework in which carriers operate—reflecting concerns for competition, environment, ATC user fees, etc.—could substantially affect air carriers' operations and might call for changes to our models. Changes in the economic forecasts that drive the TAF forecasts will not change our models, although, of course, changing these inputs will change the model forecasts. Unforeseen technological breakthroughs, or political instabilities like terrorist attacks or war, could also have substantial impacts on air carriers.

More recent work using this approach has quantified the impact of other airline strategies in response to congestion. These strategies include more point to point operations, additional hub airports, nighttime operations, and larger aircraft. These results will be included in the conference paper if permission is received from the sponsoring agency.

Author Biographies

Peter Kostiuk is the Program Manager for Technology Assessment at the Logistics Management Institute. His research in recent years has focused on quantifying the operational and economic impacts of aviation technology. Under a NASA-sponsored program, he and his colleagues at LMI recently developed the Aviation System Analysis Capability, an integrated suite of models and databases for evaluating the impacts of changes in technology, policy, and procedures on the air transportation system. He holds a B.A. from Rutgers University and received his Ph.D. in economics from the University of Chicago in 1986.

Dou Long is a Research Fellow at the Logistics Management Institute. His recent work has been focused on modeling operational changes in the National Airspace System. He received his B.E. in electrical engineering from the University of Science and Technology of China in 1984 and his Ph.D. in systems engineering from the University of Virginia in 1990.

David Lee is a Senior Research Fellow at the Logistics Management Institute, where he has worked since 1993. Previously, he worked at the US Department of Defense as the Director of the Procurement Cost Analysis Division, and served as head of the Department of mathematics and Computer Science at the USAF Institute of Technology. He was a Visiting Professor at the Von Karman Institute in Brussels. He holds a Ph.D. in Mechanics from the Illinois Institute of Technology and a M.S. in Applied Mathematics from Brown University.