MODELING THE CAPACITY AND ECONOMIC EFFECTS OF ATM TECHNOLOGY

Dr. Peter F. Kostiuk
Dr. David A. Lee

Logistics Management Institute
2000 Corporate Ridge
McLean, VA 22102−7805
E−mail: pkostiuk@lmi.org

Keywords: air traffic management, queuing models, economic analysis

1. Introduction

The large investments required to develop the next generation of tools and methods for air traffic management highlight the need to generate credible estimates of the capacity effects of those changes and their economic impacts on users. To support the NASA aeronautics research program, over the past several years the Logistics Management Institute has developed several models to evaluate the impacts of advanced aviation technologies. Much of the work has been performed as part of a comprehensive model development project called the Aviation System Analysis Capability (ASAC). ASAC is a multi−year model and database development effort designed to provide a comprehensive ability to analyze the impacts of advanced aviation technologies on the air transportation industry. ASAC has been applied to several studies of NASA R&D programs, including Terminal Area Productivity (TAP) and Advanced Air Transportation Technology (AATT).

Current models applicable to ATM include analytical models of runway and airport capacity, a queuing model of en route and terminal radar approach control (TRACON) sectors, queuing network models of the US National Airspace System, optimal route generators, and economic and cost models. The fast−time analytical runway, airport, and sector capacity models incorporate extensive feedback from controllers and capture the key effects of technology improvements and procedural changes. The network models integrate the airport capacity and sector models and provide statistics of throughput and delay. They allow estimates of the relative impacts of decision support tools for en route controllers or tower and TRACON controllers, as well as other changes in ATM concepts. The economic model generates forecasts of the effects of technology infusions on air travel demand, airline costs and profits, and aircraft fleet requirements.

Our analysis approach emphasizes a hierarchy of models that analyze components of the ATM system at different levels of detail, to support different stages of the concept exploration and development process. The models cover airspace analysis, economic benefits, and system safety analysis.
Much current research aims to increase airport throughput. To evaluate those concepts, we use models to estimate airport capacity, a system−wide queuing network model, and economic models. The next sections briefly describe these models.

2. Airport Capacity Modeling

Analytic calculations of the statistics of runway capacity were made as long ago as 1960[1]. In 1972, the FAA initiated a program to develop airport capacity and delay models, that led to the FAA Airfield Capacity

Model. Documented in several reports[2,3], the model was upgraded in 1981[4]. Ashforth and Wright[5] and Odoni[6] give comprehensive reviews of capacity models.

LMI's Runway and Airport Capacity Models were originally built by LMI to estimate the capacity and delay impacts of NASA's TAP program. The Runway Capacity Models form the building blocks of airport−specific capacity models that reflect the runway configurations and operating procedures used by controllers at individual airports. Both single and dual runway versions of the Runway Capacity Model are available. The dual runway version is used to analyze operations on closely spaced runway pairs at airports such as ATL and DFW.

The Runway Capacity models approach the problem of calculating airport capacity from the controller's perspective: given the quality of information on aircraft speed and position, wind speed, and operating procedures, what spacing should be applied to ensure that operating rules are not violated? The models incorporate the information statistically and calculate a buffer that allows the controller to meet a specified level of confidence that all restrictions (such as minimum horizontal separation and single runway occupancy) will be met. The effects of changes in technology or procedures are incorporated through changes in the parameters that describe information quality or operating procedures. Table 1 lists the key model parameters. For example, moving from radar surveillance to GPS position reporting may reduce position uncertainty from 0.25 nautical miles to 100 feet, or wake vortex prediction systems at the airport could allow for reduced minimum separation from 3.0 nautical miles to 2.5 or less.

LMI currently maintains models for 10 airports: ATL, BOS, DTW, EWR, JFK, LGA, ORD, DFW, LAX, and SFO. Given the set of runway configurations they use, the models can be modified to accommodate other airports. The models estimate hourly capacity based on weather conditions and parameters that reflect operating procedures and available technology. The model selects the runway configuration and arrival and departure mix that most closely meets the arrival/departure demand profile. We currently have 35 years of hourly weather for each of the 10 airports, and 30 years of weather for most of the other pacing airports in the US. The principal weather inputs are ceiling, visibility, wind speed, and wind direction. Given those inputs, the model estimates the capacity for all legal configurations and chooses the maximum (subject to other constraints, such as noise restrictions). The model then creates an output file with hourly capacities for the time period for which weather data were input. For most analyses, we use at least a full year of data in order to capture the variability in capacities and to avoid the ambiguities of extrapolating from a limited number of days.

Table 1. Capacity Model Parameters

LMI Runway Capacity Model	
p_i, fraction of aircraft in class I	σ_{Di}, s.d. of departure speed
S_{ij}, miles−in−trail minima	σ_x, s.d. of position uncertainty
V_i, approach speeds	σ_{Vi}, s.d. of approach speed
D, common path length	σ_w, s.d. of wind speed
R_{ai}, arrival ROT	c, mean communications

	delay
σ_{Ai}, s.d. of arrival ROT	σ_c, s.d. of communications delay
R_{di}, departure ROT	V_{di}, departure speed
σ_{Di}, s.d. of departure ROT	WAKE flag
D_d, distance−to−turn on departure	MAX_HISTORY (integer)
Note: Subscripts indicate variation with aircraft class. ROT = runway occupancy time; s.d = standard deviation.	

The airport model is an airport−unique executive controller. Each airport model consists of a set of runway configurations and the rules for their operation. A runway configuration can range from a single runway to as many as 7 runways operating simultaneously. The rules, which are obtained from interviews with airport and TRACON controllers, define the appropriate airport runway operating configurations for the various meteorological conditions (i.e., ceiling, visibility, wind speed, and wind direction). The rules also specify runway use with respect to all arrival, all departure, or mixed use, plus constraints due to dependent operations. Finally, the rules cover special noise requirements, conditions for three− and four−runway operations, and any other operational factors. The airport model calls the runway modules and combines their outputs to produce for each TAP technology level a set of arrival and departure capacity curves that are functions of meteorological conditions. In some cases, for a given meteorological condition, different configurations are used depending on the need to handle more departures or more arrivals.

We integrated an Airport Delay Model with the airport capacity models to produce hourly delay estimates for an extended time period. Airport demands, measured in arrivals and departures per hour, are applied against the capacity estimates through a queuing model that estimates delays. The cost of those delays can be estimated using cost models developed by LMI. By analyzing at least one year of weather, and more commonly 30 years, the analysis provides information not only on average or total delay reduction, but also on the distribution of delay. Consequently, we can quickly generate data on the variability of flight times and obtain insight into the possible effect of changes in predictability on airline operations.

The Runway and Airport Capacity Models have been used to estimate the benefits of the NASA TAP program. The goal of the TAP research program is to safely achieve visual flight rule (VFR) capacity in instrument flight rule (IFR) conditions. In cooperation with the Federal Aviation Administration (FAA), NASA’s approach is to develop and demonstrate airborne and ground technology and procedures to safely reduce aircraft spacing in terminal areas, enhance air traffic management and reduce controller workload, improve low−visibility landing and surface operations, and integrate aircraft and air traffic systems. By the end of the decade, integrated ground and airborne technology will safely reduce spacing inefficiencies associated with single runway operations and the required spacing for independent, multiple−runway operations conducted under instrument flight rules.

The NASA TAP program consists of four major program elements: Reduced Spacing Operations (RSO), Low Visibility Landing and Surface Operations (LVLASO), Air Traffic Management (ATM), and Aircraft/ATC System Integration. The RSO element focuses on building systems to reduce current aircraft spacing standards in terminal areas. LVLASO concentrates on developing technologies to cut delays on the runways and taxiways during periods of poor visibility. The ATM element builds on the Center TRACON Automation

System (CTAS) Program currently being supported under the NASA base program and the FAA Terminal Air Traffic Control Automation (TATCA) Program.

The fourth element of TAP, Aircraft/ATC Systems Integration, focuses on ensuring that the various systems developed under the other elements fit consistently into the overall system. The goals of this element are threefold: (1) Ensure coordination and integration between airborne and ground−side elements; (2) provide flight facility support; and (3) develop and maintain the systems focus with technology impact and cost−benefit analysis. LMI’s work has been performed as part of the Aircraft−ATC Systems Integration element[7].

At completion of TAP research and development in 2000, the technology requirements will be established by analysis and testing (validation). Hardware and software feasibility will be demonstrated by integrated tests (demonstration). The next phase of TAP development varies with the technology. Wake vortex sensors and other developmental hardware will require further engineering and manufacturing development, probably by the FAA, while software products like CTAS upgrades may need no further development. (Some modifications of software will be required to meet FAA reliability and hardening standards.) Suites of commercial off−the−shelf hardware, like flight management systems and data links, may need no further development, but will require purchase or upgrading by individual airlines. For cost analysis purposes the TAP products fall into the following categories:

- algorithms and software that can be installed in existing FAA and aircraft systems,
- validated specifications supported by feasibility demonstrations for hardware to be further developed and purchased by the FAA, and
- specifications and recommendations for new or modified commercial off−the−shelf avionics to be purchased by the FAA and aircraft owners.

Figure 1 displays the basic process for estimating the impact of TAP technologies on arrival minutes of delay. The operating meteorological conditions are IMC−2 (Category 2 and 3 IFR), IMC−1 (Category 1 IFR), VMC−2 (airport−unique reduced VFR), and VMC−1 (VFR). The levels of technology are Current Reference conditions, the 2005 Baseline, and three incremental implementations of TAP. The 2005 Baseline includes the Wide Area Augmentation System (WAAS) and the Center TRACON Automation System (CTAS). TAP 1 adds AVOSS; TAP 2 adds LVLASO reduced roll−out and turn−off (ROTO), and TAP 3 adds full ATM CTAS/FMS integration.

We focus on aircraft−minutes of arrival delay in the terminal area as the principal performance measure. Estimating delay first requires calculating airport capacities as a function of runway configurations, weather−based air traffic control operating procedures, and the TAP technology levels. Second, future demand is forecast using the predictions in the FAA Terminal Area Forecast (TAF). Next, capacity as a function weather, projected demand, and historical weather data are fed to a queuing model to generate arrival delay statistics as a function of TAP technology. The cost of a minute of delay calculated from historical airline data is used to estimate the value of the delay reductions generated by the TAP technologies. Finally, those savings are compared with the estimated lifecycle costs for the TAP systems to produce benefit to cost ratios.

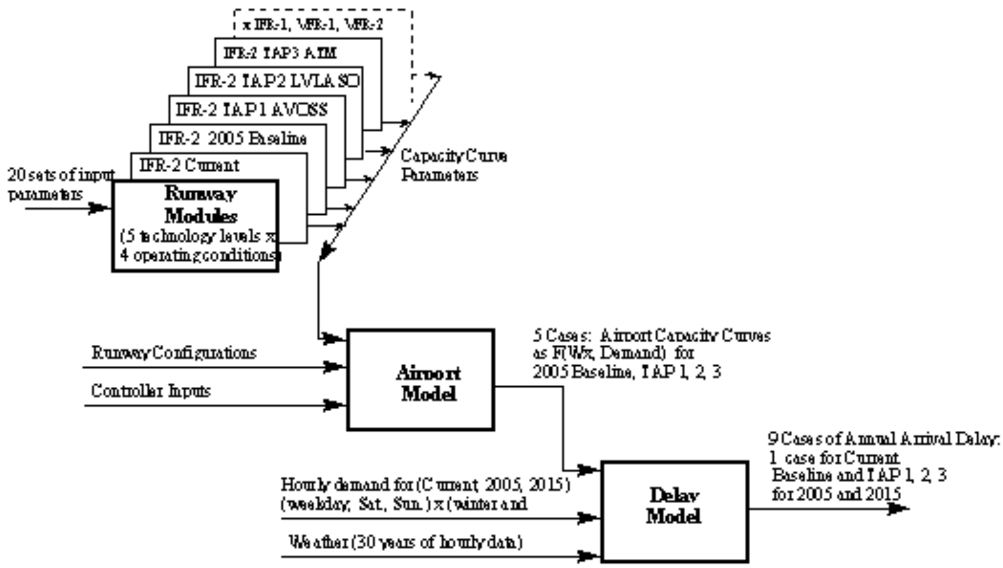

Figure 1. Airport Capacity and Delay Analysis

In the analysis, for given weather conditions at a specific airport, airport capacity is driven by the parametric variables in the capacity models. Those parameters, which include aircraft separation, approach speed, runway occupancy time, uncertainties in approach speed and uncertainties in position relate directly to controller behavior and equipment performance. TAP systems impact capacity and delay through changes they allow in the capacity model parameters.

Two steps were required to link delay reductions to changes in airline operating costs. First, we identified the elements of airline operating costs that are affected by terminal area delays. Second, we identified the relationship of those costs to the length of the delay. With the cost per minute of arrival delay thus established, it is straightforward to calculate the benefits of the TAP systems from the increases in capacity and corresponding reductions in delay.

The primary parametric impacts of the TAP technologies are:

- TAP 1 is reduced interarrival spacing, enabled by controller aids for minimum spacing between specific pairs of aircraft.
- TAP 2 is reduced interarrival spacing plus a 20% reduction in IMC−2 runway occupancy time.
- TAP 3 is reduced interarrival spacing, reduced runway occupancy time, plus additional reduced interarrival spacing enabled by real−time two−way communication between aircraft flight management systems and CTAS computers to deliver the aircraft over the threshold accurately at the time requested by the controller.

The principal analytical challenge is selecting the parameters that best reflect the impact of the technologies. For interarrival spacings affected by AVOSS, minimum spacings are reduced by about 0.5 to 1.0 nm, depending on aircraft pair. However, reductions below the 3.0 nm minimum are not possible due to ROT constraints with wet runway conditions. For large aircraft trailing large aircraft, the current baseline is 3.0 nm, under TAP 1 it is unchanged, for TAP 2 it is reduced to 2.5, and for TAP 3 it falls to 1.9 nm.

Table 2 shows the preliminary estimates of arrival delay reduction at BOS and DTW for the baseline forecasts and different technology states. Recall that the TAP benefits accrue from reduced delay. The delays are estimated for the 1993 historical demand and for projected demand levels in 2005 and 2015.

Table 2: Annual Aircraft Arrival Delay (millions of minutes)

Technology State	1993	2005	2015
BOS Current	5.5	6.8	12.2
BOS TAP 1	−	5.9	10.8
BOS TAP 2	−	4.8	8.9
BOS TAP 3	−	2.1	4.2
DTW Current	1.1	1.6	2.8
DTW TAP 1	−	1.5	2.6
DTW TAP 2	−	1.4	2.0
DTW TAP 3	−	1.1	1.4

As the table shows, the potential benefits at BOS are significantly greater than at DTW. The principal reason for this result is that DTW has sufficient IFR capacity for the next several years, due to its ability to maintain independent parallel approaches during IMC. BOS, in contrast, with key runways separated by only 1500 feet, loses one of them when ceiling and visibility drop.

One other result worth noting is the large delay reduction from TAP 2. This benefit derives from lowering wet runway occupancy times to under 50 seconds. Under existing FAA procedures, airports can reduce their IMC minimum separations to 2.5 nm when they can document ROTs less than 50 seconds. Consequently, any program that allows separations to meet that restriction generates a substantial benefit. This result also demonstrates the close dependency of the ROT and reduced separation research programs – both are required in order to achieve maximum benefit.

3. LMINET

The network model LMINET links queuing network models of airports via queuing models of TRACON and ARTCC enroute sectors. Sequences of enroute sector models, determined by an aircraft trajectory model, model flight operations between airport terminal areas. Each of these models includes several parameters, at more than one level, that may be adjusted to reflect the impacts of changes in technology or procedures. LMINET was developed under a task for AATT, and the airport models were built for the NASA TAP program.

The airport model accounts for interactions between arrivals and departures, both in the short term and in the long term. Short−term interactions are modeled by using airport capacity models that explicitly trade arrival capacity against departure capacity, to meet present demands. Long−term interactions are modeled by tracking the airport's "reservoir" of ready−to−depart aircraft; depriving an airport of arrivals for a sufficiently long time also stops departures, for lack of airplanes. Figure 2 shows a diagram of the airport model.

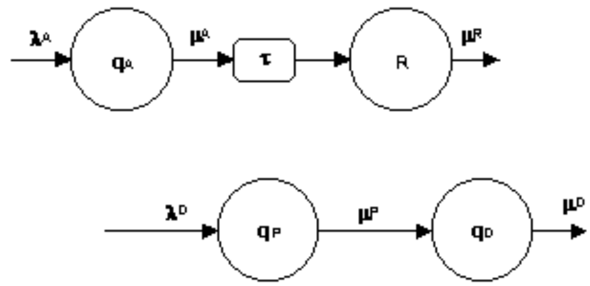

Figure 2. Diagram of the airport model

Traffic enters the arrival queue, qA, according to a Poisson arrival process with parameter A(t). Upon service by the arrival server, which is Poisson service with parameter $\mu A(t)$, and after the turnaround delay , arriving aircraft enter the ready−to−depart reservoir R. Each day’s operations begin with a certain number of aircraft in this reservoir.

In LMINET, airport arrival and departure capacities $\mu A(t)$ and $\mu D(t)$ are determined dynamically from the capacities of runway configurations. The configuration capacities are functions of weather (ceiling, visibility, wind speed and direction, and temperature) and arrival and departure demand. Runway capacities come from the LMI runway capacity model developed for NASA. Presently, LMINET models TRACON sectors, and enroute sectors, as either M(t)/D/N(t) queues or M(t)/Ek/N(t)/N(t)+q queues. The service−rate parameter D varies from sector to sector, as does the function N(t), which gives the maximum number of aircraft that a sector’s controller can accommodate simultaneously. LMINET’s outputs are statistics of arrival and departure queues with estimates of associated delays at individual airports, and statistics of delays in enroute and TRACON sectors.

Recent work has shown encouraging agreement between observations and models of ARTCC and TRACON sectors as M/Ek/N/N+q queues[8]. That is, as queues for which the interarrival times have a Poisson distribution and service times have an Erlang−k distribution; the maximum number of aircraft that can be accommodated in the sector at a given time is N; and, when N aircraft are in the sector, up to q additional aircraft may "wait" for service in the sector. (The "waiting" is accomplished by vectoring and speed changes.) Thus at a high level, enroute and TRACON sector throughput capacities are determined by the maximum number N of aircraft that the controller can handle simultaneously, and by the mean μ and dispersion parameter k of the Ek distribution. The maximum number N varies with the complexity of tasks that flights impose on the controller, with weather, and with specific features of the sectors. These features include, for example, reductions of available airspace by mountains, zones of imperfect radio and radar coverage, and the presence of more than one busy airport. A recent FAA publication implies that N ranges from 15 to 21, although some values may fall outside that interval.

One may infer values of μ and k fairly readily from observations of sector operations, for example, from ETMS data. Direct observation of N is more challenging. The sector occupancy that triggers controller actions to reduce arrivals to a given sector is not recorded, and controllers may accommodate "spikes" for short periods that would be intolerable if maintained for longer times. Surveys of controllers and supervisors may be the best source of information on N. Alternatively, the FAA maintains a list of "Monitor Alert Parameters" (MAPs) for all sectors, and these are occupancy values at which actions normally should be taken to unburden a given sector. The MAPs offer some guidance to values of N.

LMINET is adjustable at several levels, with a variety of parameters, to model effects of changes in procedures and automation. Since LMINET operates at a fairly aggregated level for the entire NAS, it provides a convenient way to estimate the possible benefits of various mixtures of technologies and procedural modifications. It also provides a consistent model for incorporating improvements in both terminal and en route operations, thereby supplying a consistent mechanism for analyzing tradeoffs among different combinations of technologies. Combined with high fidelity simulation models, such as RAMS or TAAM, it supplies a framework for evaluating ATM approaches from the early concept exploration stage through detailed operational assessments. LMINET is currently being used to estimate the potential benefits of several AATT decision support tools.

4. Flight Segment Cost Model

Many ATM analyses require aircraft trajectories, either from actual data or simulated tracks based on airline objectives, aircraft performance, and air traffic procedures. To provide such information, we built the Flight Segment Cost Model (FSCM), which generates an optimized trajectory based on aircraft performance and optimization rules. The Mission Generator module within the model is a tool to determine trajectories flown, fuel consumed, and time by phase of flight for one or more flight segments.

The ASAC Flight Segment Cost Model is hosted by LMI on a UNIX server. Remote (and simplified) access is provided by World Wide Web pages available to authorized users at http://www.asac.lmi.org.

The aircraft performance parameters used in the model come from the Eurocontrol Base of Aircraft Data (BADA) models. The model also accommodates inputs from the NASA preliminary aircraft design models ACSYNT and FLOPS, or any set of aircraft performance data in the appropriate format.

The model computes fuel burns and flight times on four−dimensional optimum paths, and for flights on FAA preferred routes and great circle routes with one step climb. Winds aloft data for each day in 1995 (except for three missing days) are available for use.

From these basic results, we obtain two fundamental indicators of free flight benefits to air carriers. The first indicator is the distribution of fuel savings, for operations on optimal routes instead of on preferred routes, for specified target stage times. This illustrates the potential for fuel savings when an airline operates on optimal routes, with schedules conditioned by preferred route flight times.

The second indicator of benefits is the distribution of flight times, when flights are operated on optimal routes with a fuel burn target that is representative of flights on preferred routes at the originally specified target time. This illustrates the potential for savings in operating costs from operating on optimal routes. More elaborate calculations, which we hope to make in extensions to this work, could illustrate the potential for operations on optimal routes to reduce fleet size.

The model calculates fuel burns and times for six phases of flight: Taxi−out, takeoff, climb, cruise, descent, and taxi−in. For most ATM analyses, we focus on the cruise phase to determine the optimal or constrained trajectories. We decomposed the trajectory optimization problem into two separate problems. First, we used the route waypoints as design variables to minimize the fuel, with Mach number and altitude profiles fixed. Then, holding the waypoints constant, we further minimized fuel burn through variations in the Mach number and cruising altitude. Finally, to be sure the waypoints were not suboptimal in the second problem, a final recheck of the path optimization problem was performed. This sequence may be viewed as applying the familiar and robust relaxation approach to the full optimization problem.

Two companion models work with the Mission Generator to comprise the FSCM. The Cost Translator converts flight resources into dollar costs. It is designed to work with the Mission generator, but it also operates independently. The Network Cost Model can be used in two ways: It can be used to summarize the costed output from the FSCM of a network of flight segments, or it can be used independently to estimate the costs of flying an airline’s network of flights where the number of departures and average stage length of flights flown by each aircraft type are known.

In a study for the NASA Ames Research Center AATT program, we estimated the savings to airlines from flying four−dimensional optimal routes versus flying the FAA preferred routes. This work applied LMINET and the FSCM described previously. We compute fuel burns and flight times for flights on four−dimensional optimum paths, and for flights on FAA preferred routes with one step climb, for all 362 days of calendar year 1995 for which we have wind data.

We analyzed fuel and time for turbojet operations between LAX and BOS, SEA and ORD, ORD and MIA, ORD and DFW, SEA and LAX, BOS and MIA, and DCA and ATL. The OAG indicates that carriers serve these routes with the aircraft that we modeled, the B757−200. These cities range over many regions of the contiguous United States. They provide a reasonably broad range of stage lengths, ranging from 2,263 nautical miles for LAX–BOS to 474 nautical miles for DCA–ATL. ORD is heavily represented, but it is heavily represented in air carrier operations, too.

In addition to results for turbojet aircraft, represented by the B757−200, we analyzed time and fuel for one route, CHT − IAD, for turboprop aircraft. We used the BADA model of the BAE 41 as a representative turboprop airplane.

Comparing the fuel burn on the FAA Preferred route from BOS to LAX that gave minimum fuel burn with fuel burn on the optimal route, for 359 days of wind aloft data, resulted in the histogram of Figure.

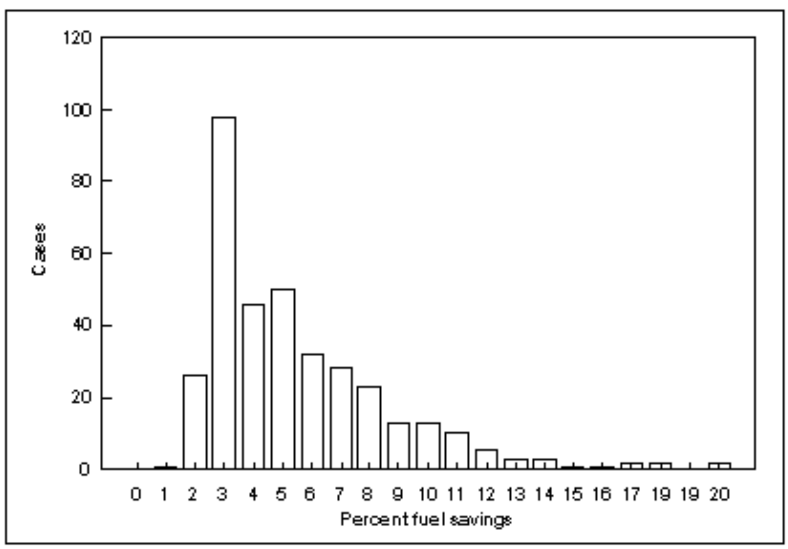

Figure 3. Histogram of Fuel Savings, BOS−LAX

The mean savings is 5 percent. The mode of the savings is approximately 3 percent. The total amount of fuel saved is 796,000 pounds. Using a fuel cost of $0.13 per pound and an average of 3.5 flights per day, this saves $362,000 a year on B757 flights from Boston to Los Angeles.

Our optimal trajectories are minimum fuel trajectories, subject to a maximum−time constraint. This results in flights at the specified maximum time, unless the absolute minimum−fuel time is less than the specified time. The maximum−time constraint that we chose for BOS–LAX flights, 5.4 hours, was active in almost all cases for flights on the preferred routes. In all but 7 days, flight time on the preferred route was within 5 minutes of 5.4 hours. On the optimal routes, however, the 5.4 hour constraint was inactive on many days, and this resulted in modest but nonetheless useful time savings in addition to fuel savings. On 55 days flights on optimal routes "saved" more than 5 minutes over the specified maximum time.

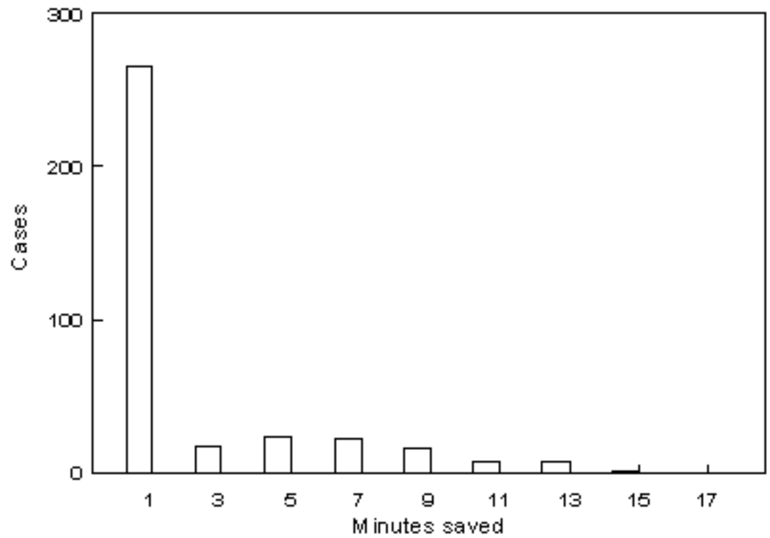

Figure 4. Histogram of Time Savings, BOS−LAX

Table 3. Summary of Results for 14 Routes			
City pair	Average fuel # savings per flight	Average fuel % savings per flight	Average minutes savings per flight
LAX−BOS	2,508	6.7	11
BOS−LAX	2,347	5.0	3
SEA−ORD	1,390	4.9	1
ORD−SEA	1,355	4.0	1
ORD−MIA	1,059	5.0	4
MIA−ORD	655	3.2	1
ORD−DFW	955	6.1	3
DFW−ORD	924	6.3	0
SEA−LAX	904	5.5	2
LAX−SEA	549	3.2	1
BOS−MIA	799	3.7	7
MIA−BOS	1,305	6.4	2
DCA−ATL	569	5.0	1

ATL−DCA	739	6.7	0

5. Air Carrier Investment Model

The increased need to demonstrate economic benefits from changes in technology and to justify government investments in air transportation highlights the need for economic models of the airline industry. To meet the demand for estimates of economic effects of advanced technologies, we use a variety of cost and economic models to translate the system impacts from the capacity and network models into financial measures. The cost models use the extensive Form 41 data on US carriers to calculate the benefits of reduced flight time and fuel consumption. Measuring the economic impacts on the industry, and the US economy, requires a more complex modeling approach. In this section, we describe the Air carrier Investment Model, which aims to forecast changes in industry costs, output, and demand attributable to changes in technology and procedures.

In creating the LMI Air Carrier Investment Model, we had some specific goals in mind. A primary objective was to generate high level estimates from broad industry−wide supply and demand factors. We envisioned being able to forecast the demand for air travel under a variety of user−defined scenarios. From these travel demand forecasts, we could estimate the derived demand for the factors of production, most importantly the number of aircraft in the fleets of U.S. scheduled passenger air carriers. We could also gauge the financial health of the airline industry as expressed in its operating profit margins.

To create the model, we identified key airports from which flights originate and developed airport−level demand models for passenger service provided by major air carriers. Furthermore, we linked the carrier−specific demand schedules to an analysis of the carriers' production functions expressed in terms of the prices of the major inputs−−labor, fuel, materials, and flight equipment. Flight equipment was modeled in an especially detailed way by incorporating some key characteristics of aircraft.

The resulting cost model generates derived demand schedules for the factors of production, in particular aircraft fleets. The derived demand schedules are functions of the price of the factor of production, prices of other factors, parameters that describe the aircraft and the network used by a carrier, and the level of passenger service supplied. A complete description of the model can be found in a recent NASA report[9].

Our first analytical task was to develop a model of demand for an airline firm's provision of passenger service. From a particular airport at origin i, carrier j will generate a certain level of passenger traffic. The U. S. Department of Transportation's Origin and Destination data record a one in ten sample of all tickets and from these, the revenue passenger mile (RPM) service originating at a particular airport for a particular carrier was constructed. Demand for a carrier's service is driven by the carrier's yield (measured by the average ticket price for flights originating at airport i divided by the average number of revenue passenger miles flown), its competitors' yields, and the size and economic prosperity of the market. We modeled the economic characteristics of the Standard Metropolitan Statistical Area (SMSA) surrounding a particular airport in terms of the area's population, per capita income, and unemployment rate. The period under

$$(1)\quad q_{t,i,j} = D_{t,i,j}(p_{t,i,j}, p_{t,i,c}, x_{t,i})$$

consideration was from the first calendar quarter of 1977 until the last calendar quarter of 1992. The demand function in equation form was given by:

where $q_{t,i,j}$ is the scheduled demand (in RPMs) originating at time t from airport i for carrier j; $p_{t,i,j}$ is the average yield for service originating at time t from airport i for carrier j, $p_{t, i, c}$ is the average yield for the other carriers generating traffic at time t from airport i, and $x_{t,i}$ are the other demand characteristics at time t for airport i. Conventional treatments for firm and airport fixed effects were used. These effects capture those important characteristics of a particular city which are not easily measured, such as tourism effects. We utilized a log−log specification for equation (1), so that the regression coefficients may be interpreted as elasticities.

Total demand for an air carrier's passenger service was then constructed by summing over the airport−specific demand equations. In terms of equation (1), the total demand for a carrier's

$$(2) \quad q_{t,j} = \sum_{i=1}^{ap} q_{t,i,j}$$

service is given by: (2) and it is the total $q_{t, j}$ that was linked to the firm−specific cost model.
The second major component of our econometric study explains total carrier costs in terms of the level of service, factor prices, network configuration, and a set of characteristics for aircraft flown. The cost analysis was based mainly on observations from the Department of Transportation Form 41 data discussed in more detail in Appendix A. From the Form 41 data, we generated a separate set of demand equations for each of the carrier's factors of production, based on standard economic assumptions concerning the cost−minimizing behavior of a carrier. These demand equations in turn permit examinations of the impact of varying factor prices and factor productivities, fleet and network configurations, and aircraft operating characteristics.

Scheduled RPM traffic for carrier j at time t was constructed as the sum of originating traffic supplied by the carrier for all airports from which it offered flights. This was the first of the two outputs considered in the cost function below. The second was the level of non−scheduled RPM service. The two generic output categories at time t for carrier j are designated $y_{t, j, 1}$ and $y_{t, j, 2}$ for scheduled and non−scheduled RPM demand, respectively. The factors of production are labor, energy, materials, and capital. In the model, capital refers to aircraft fleets only. Other (non−flying) capital such as ground structures and ground equipment are included in the materials category. Omitting the time and firm subscripts, the transcendental logarithmic (translog) cost function is given by:

$$\begin{aligned}(3) \quad \ln C = \alpha_0 + \sum_{i=1}^{2} \alpha_i \ln y_i + \sum_{i \le j}^{2} \sum_{j=1}^{2} \\ \sum_{i=1}^{4} \beta_i \ln w_i + \sum_{p \le q}^{4} \sum_{q=1}^{4} \beta_{pq} \ln w \\ \sum_{i=1}^{4} \rho_i \; \mathit{aircraft\ attributes}_i \ln w_{capital}\end{aligned}$$

In (3), w_i is the input price for factor I. The variables *aircraft attributes* and *network traits* are defined below.

Cost shares for labor, energy and materials are given by:

$$(4) \quad M_i = \beta_i + \sum_{j=1}^{4} \beta_{ij} \ln w_j$$

The cost share for aircraft capital is:

$$(5)\quad M_{capital} = \beta_{capital} + \sum_{j=1}^{4} \beta_{capital,j} \ln w_j + \sum_{j=1}^{4} \rho_j \, aircraft\ attributes_j$$

The translog cost equation can be viewed roughly as a second−order approximation of the cost function dual to a generic production function. Symmetry and linear homogeneity in input prices are imposed on the cost function by the restrictions:

$$\alpha_{ij} = \alpha_{ji}, \forall i, j;\ \beta_{ij} = \beta_{ji}, \forall i, j;\ \Sigma_i \beta_i = 1;\ \Sigma_j \beta_{ij} = 0;\ and\ \Sigma_j \alpha_{ij} = 0.$$

Summary statistics based on the translog cost equation and its associated share equations are provided by the Morishima and Allen−Uzawa substitution elasticities, and several measures of returns to scale can be obtained from the parameter estimates.

Aircraft attributes are modeled from various characteristics of the aircraft fleet. A major component of airline productivity growth is measured by changes in these attributes over time. For example, newer aircraft types−−all other things being equal−−are expected to be more productive than older types. The most significant contribution to productivity growth in the 1960's was the introduction of jet equipment. While this innovation was widely adopted, it was not universal for carriers throughout the data sample. Newer wing designs, improved avionics, and more fuel efficient propulsion technologies also make flight equipment more productive. Once an aircraft design is certified, a large portion of the technological innovation becomes fixed for its productive life.

The aircraft characteristics included in the model are average aircraft age, average size (as measured by seats), and percentage of fleet that is wide−bodied.

Two attributes of the carrier’s network are also included in the model: average stage length and passenger load factor. Stage length allows us to account for different ratios of costs due to ground−based resources compared with costs attributable to the actual stage length flown.

Shorter flights use a higher proportion of ground−based systems per passenger−mile of output than do longer flights. Also, shorter flights tend to be more circuitously routed by air traffic control and spend a lower fraction of time at an efficient altitude than longer flights. Passenger load factor can be viewed as a control for capacity utilization and macroeconomic demand shocks. It has also been interpreted in many transportation studies as a proxy for service quality.

The cost data follow sixteen domestic air carriers with quarterly observations between 1977 and the end of 1992. These firms are the set of former certificated carriers that existed throughout the study period and account for well over 95 percent of the domestic air traffic.

The joint model of supply and demand for commercial passenger air service specified in our study and the inferences about the demand for airplanes that are imbedded in our econometric results allow us to simulate the effects of emerging airframe and engine technologies through modifications to the characteristics of the planes in service. We can also simulate the growth in total system demand for passenger service and thus for factor inputs such as number of aircraft in the fleet.

We incorporated the estimated regression coefficients into an Excel spreadsheet model. Because one of the difficulties of a model specification with firm−specific effects is the appropriate value for the intercept term, we used actual airline data for the last two quarters of 1993 and the first two quarters of 1994 as the starting points from which to grow the forecasted time series of key values.

We follow several steps when evaluating scenarios. We first predict the change in RPMs based upon economic forecasts and the demand equation estimates. We next estimate airline revenues based upon forecast RPM growth and hypothesized changes in ticket prices. Airline operating costs are then estimated based on forecast RPM growth, changes in input prices, and changes in aircraft and network characteristics. We predict the aircraft inventory from airline operating costs, the capital share equation, and hypothesized changes in aircraft price and aircraft size. Finally, we compare forecasts from the LMI Air Carrier Investment Model with predicted changes in RPMs, aircraft fleet, and operating margins from other published forecasts.

To make the model reflect actual industry conditions more faithfully, three important characteristics of the industry are incorporated into the model:

1. competition among airlines that keeps operating profits at realistic levels;
2. links between airline costs and fare yields;
3. interdependency between fares and profitability.

The LMI Air Carrier Investment Model accommodates these features by building a target profit rate into the model. To meet the target profit rate, the model adjusts fare yields until the target is met. This approach builds directly into the forecast the impact of competition on the airline industry and allows the degree of competition to be set directly through the target margins. By choosing an appropriate profit rate, the user can also ensure that adequate capital is available to finance the purchase or lease of the aircraft needed to satisfy the growing demand for air travel.

As implemented in the model, separate target profit rates can be set for each of the four five−year intervals within the forecast period. Specifying four distinct periods permits the user to include changes in the economic environment during the forecast period. For example, many financial analysts today claim that airlines will not purchase additional aircraft until their balance sheets are "repaired." One way to implement this concept is to set a higher profit margin during the first five−year interval, and then set the target at a lower, historically reasonable level. Such a scenario will keep fares and profits at a higher level for five years, while reducing the derived demand for aircraft and other inputs.

The model does not impose the margin constraint in every single year. Instead, the model iterates changes in fare yield until the target margin in the final year of each interval is satisfied. Since the model uses a constant rate of fare change within each five−year interval, the operating margin does not equal the target until the final year of the period. In practice, the profit margin moves in equal increments within the interval. If the target margins are the same at the beginning and end of the five−year interval, the margin will be the same in each year.

This approach explicitly lets fare changes be set by the degree of competition and the level of costs throughout the industry. It allows for a market−based mechanism for translating cost changes into profits and fare changes. One implication of this approach is that cost−reducing technologies will primarily benefit the traveling public and not result in higher profits for the airlines over the long run. While some airlines may benefit for a short while, competition will eventually drive fares down as most airlines adopt the cost−reducing technology. This analysis is consistent with economic theory and also appears to be an accurate description of the airline industry. The relatively low profit margins reported by the airline industry demonstrate the speed with which innovations and new technologies diffuse throughout the industry. The ease of entry for new airlines with access to cheap older aircraft keeps profit margins low, and it is unlikely that this situation will change in the near future.

To demonstrate the reasonableness and utility of the model, we evaluated a set of alternative scenarios that correspond to the effects that various NASA technologies might produce. These are summarized in Table 3.

Table 3. Baseline and Scenario Forecasts

Scenario	Gross Changes in variables (%)	Annual change in RPMs (2005−2015) (%)	Annual change in aircraft (2005−2015) (%)
Baseline	N/A	4.17	2.53
A	A/C fuel = −5	4.23	2.58
B	A/C fuel = −14 A/C price = +2	4.31	2.67
C	Flight crew = −4 A/C fuel =−4 Maintenance = −4 A/C productivity = +4	4.31	2.38

6. Conclusions

The models just described, and others as well, are currently supporting the NASA aeronautics research program. Current work underway extends the analysis to analyze the effects of selected AATT decision support tools, such as the Active Final Approach Spacing Tool (AFAST), departure planning tools, and conflict probes. Those studies will both use and extend the LMINET model to accommodate more airports and improve model operating time. The FSCM is being used to generate traffic scenarios under Free Flight and other conditions, thereby providing the data needed to examine the likely changes in sector demands and associated controller workloads as traffic patterns change.

Other studies underway apply the runway and airport capacity models to the problem of estimating the capacity effects of cockpit−based independent approaches on closely−spaced runways during IMC. A related study, using methods not described in this paper, analyzes the possible safety impacts of those systems. All of the studies generate estimates of the economic effects on aircraft operators of these advanced technologies.

References

1. Blumstein, A., *Airport Capacity: Representation, Estimation, Optimization*, Cornell Aeronautical Laboratory Report TTA−1358−8−1, 1960.
2. *Techniques for Determining Airport Airside Capacity and Delay*, Report FAA−RD−74−124, 1976.
3. *Technical Report on Airport Capacity and Delay Studies*, Report FAA−RD−76−153.
4. Swedish, W. J., *Upgraded FAA Airfield Capacity Model*, FAA−EM−81−1, Volumes I and II, 1981.

5. "Airport Capacity and Configuration," Chapter 7 of *Airport Engineering*, by N. Ashford and P. H. Wright, pp. 185 et seq., Wiley, New York, 1992.
6. Odoni, A. R., *Existing and Required Modeling Capabilities for Evaluating ATM Systems and Concepts*, MIT Report , to appear.
7. Lee, David A., Peter F. Kostiuk, Robert V. Hemm, Earl R. Wingrove III, and Gerald Shapiro, *Estimating the Effects of the Terminal Area Productivity Program*, NASA Contractor Report 201682, April 1997.
8. Lee, David A., Peter F. Kostiuk, Bruce Kaplan, Marcos Escobar, and Amedeo Odoni, *Technical and Economic Analysis of Air Transportation Management Issues Related to Free Flight*, Logistics Management Institute NS501T1, February 1997.
9. Wingrove, Earl R., Jesse P. Johnson, Robin G. Sickles, and David H. Good, *The ASAC Air Carrier Investment Model (Second Generation)*, NASA Contractor Report 201678, April 1997.