



EFFECTS OF INCREASED NONSTOP ROUTING ON AIRLINE COST AND PROFIT

Terran Melconian and John-Paul Barrington Clarke

Report No. ICAT-2001-4
September 2001

***MIT International Center for Air Transportation
Department of Aeronautics and Astronautics
Massachusetts Institute of technology
Cambridge, MA 02139 USA***

EFFECTS OF INCREASED NONSTOP ROUTING ON AIRLINE COST AND PROFIT

BY

TERRAN MELCONIAN AND JOHN-PAUL BARRINGTON CLARKE

Abstract

Delays in the United States air transportation industry are increasing every year, with correspondingly increasing costs. Delays are particularly bad at hub airports, due to the extra demand placed on these connecting points. This paper addresses one approach to help alleviate this problem, that of shifting capacity from hub-and-spoke flights to nonstop flights. In order to evaluate the effects of such a change, we analyze the market share and revenue benefits of adding new nonstop flights to a market previously served only by connecting service, and examine the actual cost of delays. The MIT Extensible Air Network Simulation, developed in support of this work, is also presented. For a sample analysis for Continental Airlines, it is found that over \$550,000 per day in additional profit could be obtained by reassigning flights away from the congested hubs.

This document is based on the thesis of Terran Melconian submitted to the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology in partial fulfillment of the requirements for the degree of Master of Science in Aeronautics and Astronautics.

Acknowledgments

This work was supported under grant [insert grant number]. The authors would like to acknowledge Amedeo Odoni for guidance in the development of MEANS, Peter Belobaba for help with airline economics, and John Hansman for insight about how the airspace system works.

Contents

1	Introduction	15
1.1	Motivation	15
1.2	Organization	16
2	Overview	19
2.1	Historical Network Development - Hub and Spoke	19
2.1.1	Supply and Demand	20
2.1.2	Frequency and Market Share	20
2.2	The Costs of Hub and Spoke	21
2.3	The Return to Nonstop Flights	22
2.4	Requirements for Analysis	22
2.4.1	Delay Propagation	23
2.4.2	Analysis Time Window	23
2.4.3	Cost and Revenue Analysis	24
3	Delay Cost Analysis	25
3.1	Prior Work in Delay and Cost Assessment	26
3.2	Fixed Hourly Cost Modeling	26
3.2.1	Form 41 Costs	27
3.3	Nonlinear Crew Costs	29
3.3.1	Restrictions on Pilot Schedule	30
3.3.2	Examination of Effects on Schedule	33
3.4	Comparisons Between Linear and Nonlinear Model	35

3.5	Conclusions	35
4	New Nonstop Analysis	37
4.1	Models	37
4.2	Methodology	38
4.3	Data Requirements	38
4.4	Market Selection	39
4.5	Frequency Calculation	40
4.6	Coefficient Calculation	40
4.6.1	Linear Solution	40
4.6.2	Nonlinear Solution	41
4.7	Results	42
4.7.1	Linear Results	42
4.7.2	Nonlinear Results	42
5	The MIT Extensible Air Network Simulation	45
5.1	Data Sources	45
5.1.1	Schedule of Flights	45
5.1.2	Airborne Performance	46
5.1.3	Ground Operations	46
5.1.4	Weather	46
5.1.5	Airport Arrival and Departure Rates	47
5.2	Calculation Methods	47
5.2.1	Flight Data Structure	47
5.2.2	Schedule Loading and Generation	48
5.2.3	Taxi and Airborne	49
5.2.4	Tower Controller	49
5.2.5	The FAA and Ground Delay Programs	50
5.2.6	The Airline	52
5.3	Results	52

6	Approach	57
6.1	Cost Calculations	57
6.1.1	Determination of Hourly Costs	57
6.1.2	Determination of Aircraft Type	57
6.1.3	Preliminary Cost Calculations	58
6.1.4	Final Cost Calculations	59
6.2	Revenue Calculations	59
6.2.1	Frequency Calculation	59
6.2.2	Preliminary Revenue Calculations	60
6.2.3	Final Revenue Calculations	61
6.3	Scenario Generation	61
7	Results	63
7.1	Preliminary Results	63
7.2	Final Results	65
7.3	Conclusions	65
8	Directions for Future Work	67
8.1	Other Airlines	67
8.2	MEANS additions	67
8.3	More Sophisticated Scenarios	68
8.4	Enhanced Calculations	68
8.5	Shared Hubs	68
A	Formulae, Values, and Data	69
A.1	Fits to Historical Data	69
A.1.1	Actual Flight Time	69
A.2	Detailed Continental Results	70

List of Figures

2-1	Simple Nonstop and Hub Networks	21
3-1	Percentage of Delays Causing Crew Disruption	34
3-2	Sample Delay-Induced Costs for Boeing 767	36
5-1	MEANS Flight Progression	48
5-2	MEANS Queue Sizes, Boston, January 28, 1999	53
5-3	MEANS Arrival Times, Boston, January 28, 1999	54
5-4	MEANS Results, BOS, IAH, and PHX, July 30, 1999	55
6-1	Information Flow	58
A-1	Airborne Time Calculated from Distance	70

List of Tables

1.1	Change in Delays for Hub and Other Flights	15
3.1	Cost Categories Included in Delay Cost Estimations	28
3.2	Historical Hourly Costs for Assorted Aircraft Types	29
3.3	Sample NorthWest Hourly Crew Rates	35
4.1	Connecting Flight Value Linear Fit Results	42
4.2	Connecting Flight Value Nonlinear Fit Results	43
7.1	Average Results of Preliminary Calculation for Continental	64
7.2	Comparison of Delays in Original and Improved Case for Continental	65
7.3	Final Calculation Results	66
A.1	Replaced Flights for Continental	71
A.2	Preliminary Calculation Results by Flight	72
A.3	Airport Abbreviations	73

Chapter 1

Introduction

1.1 Motivation

Delays in the United States air transportation industry are increasing every year. In 1995, 19% of the domestic US flights of the 10 Major airlines were delayed according to the FAA's criterion of arriving 15 minutes late or more. By 1999, this figure had risen to 26%. The situation at congested hub airports is even worse. In 1995, 18% of Continental's domestic flights were delayed, and 20% of flights through its Newark hub were delayed - levels comparable to the nationwide average. By 1999, 35% of Continental's flights through the extremely congested Newark hub were delayed, while only 26% of Continental's other flights were delayed. Table 1.1 shows a comparison for those major airlines with enough of both hub and nonstop flights for a comparison to be meaningful.

Table 1.1: Change in Delays for Hub and Other Flights

Carrier	1995		1999	
	Non-Hub Delays	Hub Delays	Non-Hub Delays	Hub Delays
AA	19.2%	20.0%	24.0%	29.9%
AS	26.8%	32.5%	20.2%	30.0%
CO	18.0%	20.3%	26.6%	34.9%
DL	18.7%	24.3%	18.6%	29.6%
UA	17.9%	24.2%	24.3%	30.7%
US	16.7%	18.5%	28.9%	38.6%

No “silver bullet” for this problem has presented itself, and none appears to be waiting in the wings. However, several approaches have the promise of incrementally ameliorating delays, and if enough of these approach are combined, they could significantly improve the situation.

This paper addresses one such approach, that of shifting capacity from hub-and-spoke flights to nonstop flights. Fundamentally, the improvement is achieved by reducing demand at hubs, where demand already exceeds capacity, and relocating this demand to less-used airports where excess capacity is available.

1.2 Organization

Chapter 2 presents an overview of the key concepts. This includes a brief explanation of the history leading to the current situation, as well as an introduction to key requirements for the analysis and decisions made for the modelling.

Chapter 3 examines delay costs in detail, and provides a framework for their calculation which provides reasonable fidelity while minimizing computational difficulty and data requirements.

Chapter 4 presents economic models for the calculation of revenue from nonstop flights in a market otherwise served only by connecting flights. This is an extension of previous models which were able to treat only those markets where all significant competition took the form of nonstop flights.

Chapter 5 describes MEANS, the MIT Extensible Air Network Simulation. MEANS was developed in support of this work, and provides detailed information about simulated flight information which is critical for delay estimation.

Chapter 6 presents the overall approach. This includes the synthesis of the work explained in the previous chapters, and methods related to the overall analysis, such as the selection of connecting flights to eliminate and nonstop flights to introduce.

Chapter 7 describes the results of this analysis. Those flights which could be improved by being switched from hub service to nonstop service are identified, and the expected economic results of such a switch are presented.

Chapter 8 suggests possible directions for future work, based on data availability, computational complexity, and planned improvements to the MEANS model.

Chapter 2

Overview

2.1 Historical Network Development - Hub and Spoke

Airline networks changed dramatically in the United States in the early 1980s, in the wake of removal of regulations from the airline industry in 1978. Under the old system, flights would take passengers directly from one major city to another, with smaller cities served by connections from a nearby major city. In the current system, most passengers reach their destination by taking one flight from their origin city to a “hub”, where many flights meet. From this hub, they then take a second flight to their destination.

The reasons for this change in network topology are numerous, but all are fundamentally related to the idea of economies of scale. It is important to understand that under the regulated system, airline fares were set in part based on airline expenses - if an airline showed high expenses, it would be allowed to raise its fares. Thus, the incentive to reduce costs was not what it would have been in a free-market situation.

The simplest manifestation of these economies of scale appear in areas such as maintenance, where staff and inventory costs can be reduced by having one central maintenance facility. A subtler and more important reason for hub and spoke arrangements, however, has to do with economies of scale applied to frequency and passenger preference.

2.1.1 Supply and Demand

Demand in the air transportation market takes the form of transportation between one city and another. However, the product which is actually supplied takes the form of seats on flight legs, which may not directly connect the origin and destination of interest to customers. Because of this, the organization of a network (point to point or hub and spoke) can have a significant effect on the relationship between supply and demand observed by the airline.

2.1.2 Frequency and Market Share

An airline product is generally viewed as a round-trip ticket from an origin to a destination; non-round-trip tickets are sufficiently uncommon that many analyses ignore them. Every airline offering service between two cities is competing in that market; the fraction of passengers travelling between the cities who purchase tickets from a certain airline is the market share of the airline in question.

Market share is largely a function of flight frequency. Thus, a carrier with ten flights per day between two cities will have significantly more passengers than a carrier with only three flights per day. This makes it desirable to have multiple flights with smaller aircraft instead of one flight with a larger aircraft. However, there is a lower bound on the size of aircraft which can be economically operated, and thus for a certain demand there is an upper bound on the number of flights which can be offered.

A hub and spoke network topology has a consolidating effect which makes it possible to justify more flights to each city. Suppose that the traffic between cities A and B, A and C, and A and X is enough to justify two flights each per day with 100-seat aircraft. Suppose we introduce a hub at X, such that passengers going from A to B go via X. However, passengers from A to C will also be going via X, and of course so will the ones going straight to X. This is shown in Figure 2-1. Thus, instead of having flights with 100-seat aircraft from A to B, we have flights from A to X. The number of flights will be the same, namely six. However, instead of having two flights to B, two to C, and two to X, there will simply be six identical flights to X. Thus, the passenger at A who wants to go to B will no

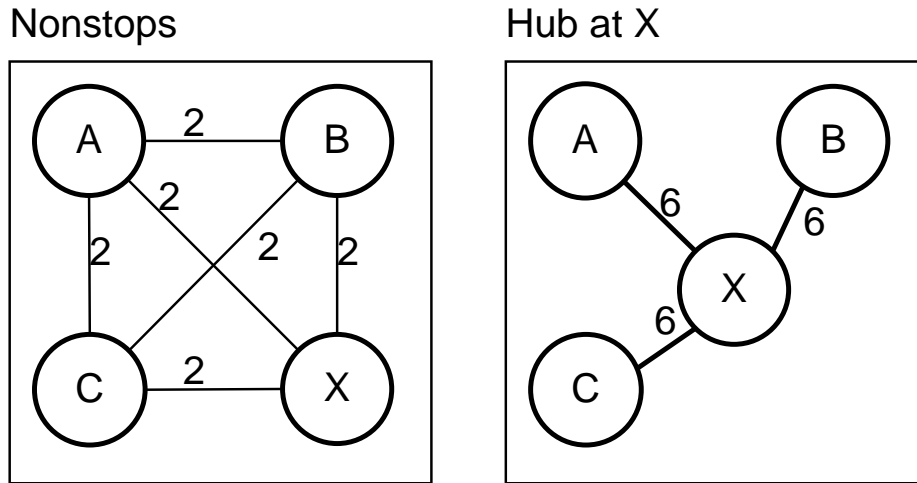


Figure 2-1: Simple Nonstop and Hub Networks

longer see two flights, he'll see six. Exactly the same happens on the second flight segment, from X to the destination city. This increased frequency provides a competitive advantage to the carrier with the hub-based flights, which realizes a higher market share and increased revenue. Alternatively, the carrier could use larger aircraft and reduce the number of flights from each city. This would still provide higher frequencies than the initial nonstop system, and it would also reduce the total number of legs, cutting costs.

2.2 The Costs of Hub and Spoke

The advantage of hub and spoke networks is that they concentrate traffic. The disadvantage of hub and spoke networks is that they concentrate traffic. Consider the example from above. Suppose in the non-hub case, there was enough traffic from A to X to justify two flights per day, and the same for B-X and C-X. This is a total of six flights per day at X. Now, suppose we start routing all traffic through X. We now have (as described above) six flights from A to X, six flights from B to X, and six flights from C to X. The traffic at X, our hub, has tripled!

This becomes a problem when the airport at X starts running out of capacity. Today, airports are running out of capacity, and hub airports are running out first due to the con-

centrating effect described above.

For example, in 1999, there were 14,900,000 enplaned passengers at Newark and 12,500,000 at Orlando. However, because Newark is a hub, there are many operations which are not reflected in the passengers counts, as they refer to connecting passengers - Newark had 457,000 operations, whereas Orlando had 366,000 operations. Furthermore, Orlando had only 0.6% delayed flights, but Newark had over ten times as many.[1]

2.3 The Return to Nonstop Flights

In a world of uncongested airports, the traffic-concentrating effect of the hub design is highly desirable. It would be foolish to remove a flight through a hub in order to add a nonstop flight between two non-hub cities, as this would reduce all of the advantages that made the hub desirable in the first place.

However, in the real world of very congested airports, congestion means delays, and delays cost money. While it might not be desirable to replace a ideal hub flight with an ideal nonstop flight, it can be very desirable to replace a delayed and costly hub flight with a nondelayed nonstop flight. The costs of delay can outweigh the advantages provided by routing the flight through the hub.

The purpose of this work is to motivate just such a change by analyzing its cost and revenue effects.

2.4 Requirements for Analysis

An analysis of the differences in revenue and cost between hub and nonstop flights does not necessarily require a full cost and revenue analysis of both situations. However, if a full analysis is not performed, revenue and cost must be assessed for at least those flights which change between the two situations, and the effect of the changes on the remainder of the network must also be examined.

Thus, the choice between a cost/revenue evaluation of the entire network in both situations or of only the differences depends on the extent of the coupling between the

added/removed flights and the rest of the network.

2.4.1 Delay Propagation

One of the difficulties with delay analysis and simulation is that delay from a delayed flight propagates to subsequent flights. Each flight, in order to depart, requires an airframe, flight crew, and cabin crew. In the worst case, the airframe, flight crew, and cabin crew from an incoming flight will be assigned to three different subsequent flights. One incoming flight can therefore delay three outgoing flights. If crews are not assigned as a group, this becomes even worse. [8]

Because historical crew schedules cannot be obtained from airlines, this work considers only the propagated delays from airframe, assuming that reserve crews are available to staff the aircraft should the regularly scheduled crews be unavailable.

2.4.2 Analysis Time Window

The majority of air transportation occurs during the daytime. Thus, excess capacity is almost always available at night, even at major hubs. It is therefore almost always the case that any delays due to reduced capacity can be made up at night, allowing the next day to start on schedule. Thus, analysis of a single day is adequate to capture all direct effects of propagated delays; i.e. delays from one day will not continue into the next day.

There are additional effects, however, which cannot be captured with a single day's worth of information. If a flight is cancelled, the aircraft which was to have flown will start the next day in a different place. The aircraft must be repositioned to where it's supposed to be before its first scheduled flight, or another aircraft must be substituted. There are costs associated with both options, and these costs can be examined only by looking at a period significantly longer than one day.

Algorithms for optimizing this recovery are themselves currently a significant area of research. In order to focus on the hub vs. nonstop tradeoff, the analysis herein has limited itself to time windows of one day, to avoid influencing the results by better or worse methods of long-term disruption recovery.

2.4.3 Cost and Revenue Analysis

Because of the propagated nature of delays, it is extremely difficult to isolate the cost of delays due to an individual flight. The cost of delays are therefore analyzed on a network level. The MEAN simulation, explained in detail in chapter 5, provides the tools required to examine the delays over the entire network.

Revenue is also difficult to isolate. Consider the simple case of a linear relationship between frequency and market share. If airline Q has two flights from A to B each day, and airline R has one flight, airline Q will get $\frac{2}{3}$ of the passengers. However, if airline Q cancels one of its flights, it will not lose half of its passengers. Assuming the flights are not full, airline Q will get half of the total passengers with one flight. This is a loss of only 25% of its previous revenue, not half. As a result of this effect, passenger revenue (a linear function of the number of passengers) is also calculated at a network level.

Chapter 3

Delay Cost Analysis

It has long been known that it is not possible to calculate the revenue from or cost of a single flight leg. In the case of revenue, this is because demand is generally measured in terms of an Origin-Destination market; any flight leg will carry passengers from many origins going to many destinations, and it is not at all clear how much of the revenue from the passenger in question should be allocated to the particular leg. In the case of costs, there are numerous costs associated with the network which are difficult to assign to a single leg. For example, the cost of crew is highly dependent on the number of hours that they are able to work in a month, and this is in turn highly dependent on the details of the schedule. It is extremely difficult to assign these costs to individual flights.

The costs of delays come from lost revenues and additional costs; as we already have difficulties assigning revenues and costs to flight legs, it is clear that we will also have difficulty determining the cost of delays from a single leg. If, for example, a flight is delayed long enough that certain passengers miss their connection, what is the additional cost or lost revenue from this? These passengers must be accommodated on other flights, either by finding space on other flights or buying them tickets with another carrier. The former is usually less costly than the latter, but the extent to which passengers can be accommodated on other flights is dependent on when the next flights between the cities in question are scheduled to occur, and how many seats are available. Alternatively, consider the delay of one flight due to the late arrival of the aircraft required to fly it. In the worst case, one late-arriving flight can delay three more, if the cockpit crew, cabin crew, and

aircraft are to be used in three different outgoing flights and no replacements or spares are available. However, the extent to which this occurs depends entirely on the timing of subsequent flights and the resources available at the airport at which the flight is late to arrive. Thus, we see that the cost of delay on a single flight is dependent not only on that flight but on other flights as well. This dependence is at the heart of the difficulty in modeling and estimating the cost of delays.

3.1 Prior Work in Delay and Cost Assessment

A great deal of work is currently being done to model and improve the delays in the National Air System. One of the most interesting of these is the Approximate Network Delay model, which attempts to model the propagation of delay throughout a network, thereby capturing part of the dependence that makes delay modeling difficult.[8] However, this model does not assign dollar values to the costs, but simply reports delays as times.

Those models which assign a dollar value to the cost of delay tend to use a fixed amount per hour, possibly depending on the location of the aircraft during the delay (air, ground, etc.)[7] These values do not take into account the differing costs of different aircraft. Some work is currently being done using the Department of Transportation Form 41 data, which provides information on the cost of aircraft depending on the aircraft type, but no work using this methodology appears to have been published yet.

3.2 Fixed Hourly Cost Modeling

Fixed hourly costs are the current state of the art in modeling costs. In addition to their simplicity of calculation, there are advantages to using a linear cost representation stemming from the fact that a linear objective function makes certain optimizations more tractable. The two sorts of fixed hourly costs which are used are an aircraft-independent cost and an aircraft-dependent cost. The latter is more accurate and therefore of more interest.

3.2.1 Form 41 Costs

“Form 41” is a form on which United States airlines report operating costs to the Department of Transportation on a yearly basis. This information is broken down by category and made available in a database. The section of costs of the most interest is the Flight Operating Cost section. The costs are reported in the Direct Flying Operations (containing Pilots and Fuel), Maintenance (containing Direct Airframe, Direct Engine, and Maintenance Burden), and Equipment Ownership (containing Depreciation, Rentals, and Insurance). All of these costs are reported on a per-block-hour basis; a block hour is the time from when the wheel blocks are removed from the aircraft prior to its departure to when they are replaced after its arrival. Some of these costs are truly incurred on an hourly basis. Fuel, for example, is clearly a cost which is directly related to the number of hours that the aircraft is in use. Others, such as ownership costs, are not dependent on the number of hours the aircraft in use, and it is somewhat artificial to report them on a block-hour basis.[9]

Depending on the type of delay in question, the categories of costs which should be included in its estimation differ. If a delay occurs regularly, and is essentially permanent, then it seems most reasonable to consider this delay as a simple lengthening of the amount of time required for a particular flight. In this case, it is reasonable to include the full hourly cost reported on the Form 41 data as the cost of the delay - as permanent delays occur, it will likely be necessary to qualitatively change the schedule, and additional aircraft may be required to serve the same cities with the same frequency. In the case of a one-time temporary delay, however, it does not make sense to include the ownership costs of the aircraft in the cost of the delay - no extra aircraft are required; if the aircraft is delayed by an hour, it is simply used for an hour longer on that day. Table 3.1 summarizes the categories in which costs are reported, and indicates to which type of delays they are applicable.

Some of the classifications shown in Table 3.1 are not entirely clear-cut and are worthy of note. The direct maintenance cost is partly dependent on hours flown (for the 100-hour checks) and partly independent (the calendar-time based checks). It is shown as being included in the cost of an airborne delay, but it could be argued that it should not be included in its entirety. Fuel is not shown as being part of the cost of a ground delay, although the

Table 3.1: Cost Categories Included in Delay Cost Estimations

Cost Category	Delay Type	Permanent	One-Time Airborne	One-Time Ground
Direct Flying Operations				
	Pilots	X	X	X
	Fuel	X	X	
Flight Maintenance				
	Direct	X	X	
	Burden	X		
Equipment Ownership				
	Depreciation	X		
	Rentals	X		
	Insurance	X		

aircraft is using some fuel while it is idling on the ground. It seems that the amount of fuel used while on the ground is sufficiently smaller than to the amount of fuel used while airborne that it can be neglected, but this classification could be argued. Finally, it may be worthy of note that direct maintenance costs are not included for a ground delay, but crew costs are. This is because crew hours are calculated based on block time, whereas aircraft duty hours are calculated based on airborne time. Therefore, a delay on the ground counts as time for the pilots but not for the airframe.

Another item of note is the lack of cabin crew costs, which are included in the overall cost data reported, but not in the data broken down by aircraft type. Because the cabin crew on an aircraft typically costs significantly less than the cockpit crew, historical practice has been to ignore cabin crew costs.

Some example costs for common aircraft types, using the categorizations presented in Table 3.1, are given in Table 3.2. It can be seen from this that the cost of a one-time delay on the ground is a fraction of the total cost of operating an aircraft of the same type. By using the total hourly cost from the Form 41 data to model the cost of delays, we may be significantly overestimating them.

Of course, none of these costs include the costs of additional delays “down the line” which may be caused by the initial delay. These costs can be estimated by using a model

Table 3.2: Historical Hourly Costs for Assorted Aircraft Types

	Permanent	One-Time Airborne	One-Time Ground
A300-600	\$3,960	\$2,372	\$675
A320-200	\$1,959	\$1,798	\$579
B727-200	\$2,238	\$1,767	\$714
B737-1/200	\$1,716	\$1,215	\$486
B737-300	\$1,874	\$1,159	\$439
B737-400	\$2,041	\$1,440	\$562
B737-500	\$1,587	\$1,053	\$401
B747-100	\$6,032	\$4,323	\$1,178
B747-400	\$5,992	\$3,728	\$1,245
B757-200	\$2,316	\$2,095	\$627
B767-200	\$2,754	\$1,821	\$677
B767-300	\$3,212	\$1,971	\$723
B777-200	\$3,594	\$2,072	\$790
DC-10-10	\$4,496	\$2,944	\$882
DC-10-30	\$4,410	\$3,173	\$805
DC-9-30	\$1,654	\$1,287	\$535
F100	\$1,685	\$1,143	\$606
L-1011-1/200	\$3,970	\$3,131	\$1,045
MD-11	\$4,153	\$2,804	\$941
MD-80	\$1,835	\$1,215	\$513

such as AND to model delay propagation in conjunction with an hourly cost model.

3.3 Nonlinear Crew Costs

The hourly crew costs reported in the Form 41 data are simply the total amount spent on crew on a yearly basis divided by the number of hours that the crew worked. This includes items such as training expenses, and paying the crew for their minimum guaranteed times even if they haven't worked them. These costs may not be applicable in the case of delays, as the pilots aren't getting any extra vacation time, for example, as a result of working an extra half an hour. On the other hand, it may also be the case that the pilots are delayed so long that they are not able to fly their original schedule, and thus must be replaced. The costs of this will be higher than an hourly cost multiplied by the pilots' time.

3.3.1 Restrictions on Pilot Schedule

The restrictions on the schedule a pilot is able to work come from the FAA's Federal Aviation Regulations (FAR) and from the pilots' contract with their airline. There are three types of limitations on the crew's schedule, any of which could be violated by delays, forcing the crew to reschedule. These limitations are:

1. Restrictions on the number of hours that a crew can fly between breaks
2. Restrictions on the amount of time that a crew can be on-duty, regardless of whether this time is spent flying or waiting
3. Restrictions on the length of break that a crew must have in order to be able to return to duty

These restrictions are discussed in more detail in the following sections, with details from the FARs and from the NorthWest/ALPA contract, examined as an example.[3, 5] The focus is on the cockpit crew, as has been done in much of the literature, due to a lack of information about cabin crew costs and scheduling.[6]

Restrictions on Hours Flown

Most of the restrictions on the number of hours that a crew can fly come from section 121.471 of the Federal Aviation Regulations. The FARs require that a pilot not fly:

1. More than 1,000 hours per year
2. More than 100 hours per month
3. More than 30 hours in any 7 consecutive days
4. More than 8 hours between rest periods

It is the last of these which is of the most interest here. If the extra flight time introduced by a delay would cause a crew to exceed their monthly allowed time, for example, they can still complete their current duty period, and be rescheduled to reduce their flying

time for later in the month. However, the requirement that a crew not fly for more than 8 hours without rest is more immediate, and is the requirement most likely to force a crew to abandon their original schedule.

The wording of the regulation is significant. It reads “No certificate holder... may schedule any flight crewmember and no flight crewmember may accept an assignment for flight time... if that crewmember’s total flight time in all commercial flying will exceed [conditions].” The significance of this is that it restricts assignment and acceptance of assignment, not actual flight. Therefore, once a flight has begun, the crew need not abort it and land at an alternate airport, even if finishing the flight will cause them to exceed eight hours of flight time.

A crew will usually, in domestic operations, be scheduled to fly more than one flight during a duty period. As indicated above, if the last leg of this series is delayed, there is no effect beyond paying the crew extra for the extra time they’ve flown, and possibly adjusting their schedule for subsequent days. However, if any of the previous legs are delayed, it may not be possible for the crew to fly the last leg, as at the time of departure the expected length of the flight would put them over the eight-hour limit.

Restrictions on Break Length

The FARs also provide requirements on the length of a break which must be given in order for a crew to fly more than 8 hours in any 24-hour period. The rests must be:

1. 9 hours of rest for less than 8 hours of flight
2. 10 hours of rest for 8 to 9 hours of flight
3. 11 hours of rest for more than 9 hours of flight

However, these rest periods can be further reduced (to 8, 8, and 9 hours, respectively) if longer breaks, of lengths specified by an even more complicated set of rules, are given for the next rest.

If any delays are encountered during a series of flights, the crew will arrive later than anticipated and begin their rest later than anticipated. If the amount of rest available to them

is sufficiently shortened, it may be impossible for them to depart on their first scheduled flight the next day, as the reduced amount of rest they obtained as a result of the delay no longer meets the minimum requirements for them to fly again.

Additional requirements on rest period length can come from the airline/pilot contract. In the NorthWest case, it is specified that all breaks shall be at least 9 hours (not 8), and after more than 8.5 hours of actual flight time, the pilots must receive at least a 12-hour rest. In this case, it is possible that any of the flights in a series might be delayed sufficiently to increase the total flying time from under 8 hours to above 8.5 hours - in which case the total rest time required would actually be increased, in addition to the actual time received being reduced. This could result in the crew being unable to fly their planned departing flight on the next day.

Restrictions on On-Duty Time

The FARs do not provide any restriction on the total amount of time that a crew can be on-duty, but contracts between individual airlines and pilot unions do. The NorthWest contract, which was examined as a sample, provides the following basic restrictions:

1. No duty scheduled to start between 0501 and 2159 (local time) may be scheduled to last longer than 13 hours.
2. No duty scheduled to start between 2200 and 0500 may be scheduled to last longer than 12 hours.
3. No duty may actually last longer than the scheduled maximum plus one hour.

There are further complications and refinements. If a 6-hour break is given during a duty period which starts within certain hours, the length of the duty period can be extended to 14 hours. If a break of less than 10 hours is given before a duty period, the maximum scheduled length of the duty period is reduced to 11 hours. Further specifications exist for redevies, international flights, and other cases.

3.3.2 Examination of Effects on Schedule

Sample Schedule

Ideally, an actual historical crew schedule would have been used to examine the effects of delays at various points on the ability of the crew to complete their planned flights without violating any of the requirements on their working and rest times. Regrettably, it did not prove possible to obtain such a schedule, as these schedules are considered valuable competitive information by the airlines and are not available for research purposes. It was therefore necessary to use a hypothetical schedule instead.

To examine the costs of a crew's inability to complete their scheduled set of flights, a hypothetical schedule was therefore generated manually using a simplified set of constraints. The constraints used were:

1. No scheduled flight times longer than 8 hours per duty period
2. No duty periods longer than 13 hours scheduled or 14 hours actual
3. No rests shorter than 9 hours, scheduled or actual

These constraints obviously do not capture the entirety of the problem. It seems likely, however, that they will provide a schedule which is similar in nature to that actually used, and which will exhibit the same behavior in response to delays. The actual schedule generated was for a historical schedule with 7 aircraft and 40 flight legs. This schedule is believed to be large enough to give a representative estimation of the effect of delays while remaining a tractable problem for manual crew scheduling.

Results of Delay in Sample Schedule

The effects of delay on the schedule were examined by looking at each leg of each flight and determining the largest delay which could be absorbed without forcing a replacement crew to be used, assuming that all other legs run at their planned duration.

Three types of delays which could prevent a crew from completing their originally scheduled flights were examined; these are an increase in flight time, an increase in duty

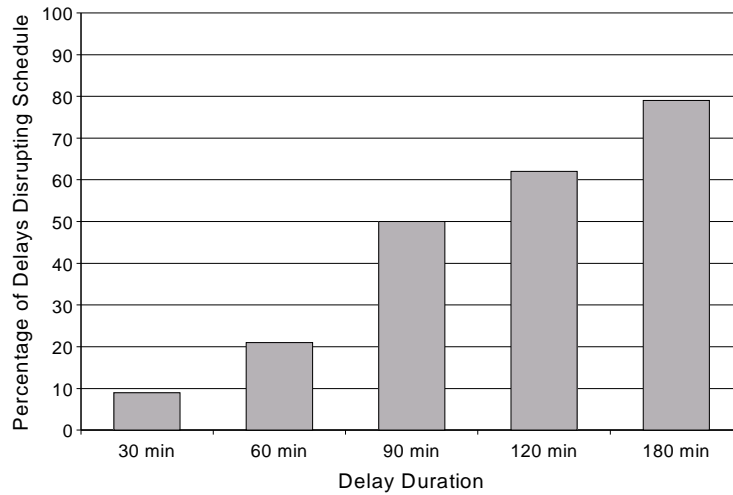


Figure 3-1: Percentage of Delays Causing Crew Disruption

time, and a decrease in rest. All three of these types appeared in the sample schedule. The majority of inability to complete the planned schedule resulted from expected flight time exceeding 8 hours due to delay in one of the legs which was not the last one.

Using this information, it was possible to calculate the likelihood that a delay of a certain length would result in a required schedule change. These results assume that delays are equally likely at all airports in the schedule. With a larger schedule, it would have been useful to examine the delay statistics of each of the individual airports, but given the size of the sample schedule available, these results would not have been statistically useful.

The results, assuming delays are equally likely everywhere, are presented in Figure 3-1.

Costs of Schedule Disruption

The actual cost of a crew being unable to complete their originally scheduled flights is of course highly dependent on how the disruption is handled. The details of how this is done are generally not available, for the reasons discussed in 3.2.1, requiring some assumptions to be made. Here, it is assumed that the crew which was not able to fly their original flights will be deadheaded (flown as passengers) back to their base, while another crew is sent out to take over for them.

The NorthWest contract indicates that pilots being deadheaded are paid at their full rate,

Table 3.3: Sample NorthWest Hourly Crew Rates

	Captain	First Officer
A320	\$175	\$124
B747-400	\$227	\$154
B757	\$183	\$124
MD-80	\$169	\$115

as if they were flying the aircraft. The cost of deadheading four crew members (two back to base and two from the base to where they must take over) is therefore quite significant. Some hourly crew rates for full-seniority crew are given in Table 3.3.

3.4 Comparisons Between Linear and Nonlinear Model

In order to compare the cost estimates provided by the linear cost-per-hour model with those provided by the nonlinear crew-schedule-disruption model presented above, we look at the average cost of a delay in the latter. The probability of a delay resulting in schedule disruption is multiplied by an average schedule disruption cost, taken from calculating the average stage length of the schedule in question. This is added to a linear hourly cost for the captain and first officer, using the rates indicated in Table 3.3.

Figure 3-2 presents a comparison between the model presented above and the fixed-hourly-cost model using only crew costs, for a 767.

3.5 Conclusions

As can be seen from Figure 3-2, the agreement between the nonlinear schedule-disruption model and the fixed hourly cost model using Form 41 data agree surprisingly well for crew costs. It may be reasonable to conclude, therefore, that the fixed-hourly-cost model is suitable for the majority of purposes, and that the nonlinear schedule-disruption model has little to offer without specific information about the actual crew schedule in place at the time of a particular delay. Such a model would be most appropriate for use at the airlines themselves, where detailed information about the actual crew schedules is available.

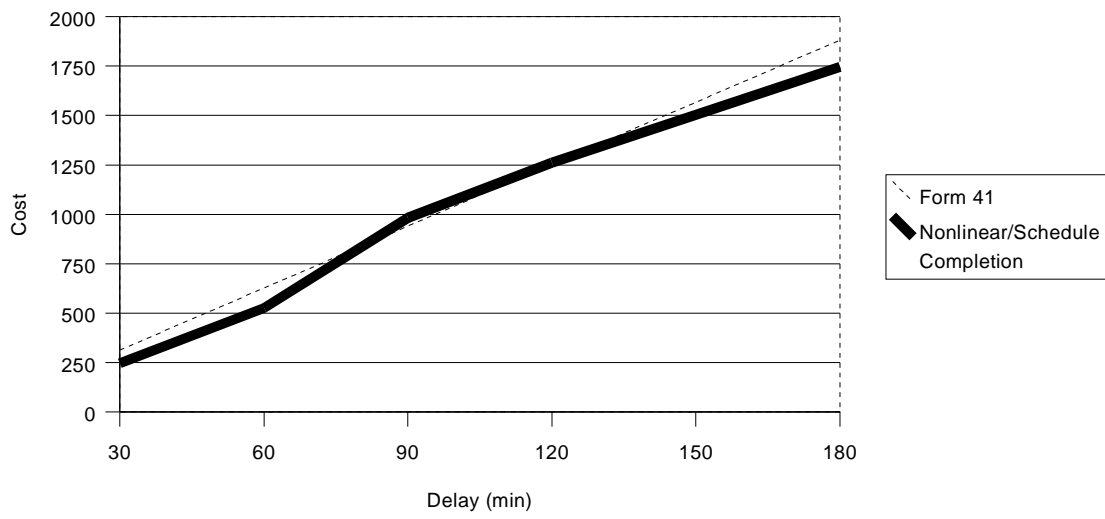


Figure 3-2: Sample Delay-Induced Costs for Boeing 767

Chapter 4

New Nonstop Analysis

When the expected profitability of a new nonstop flight in a market not currently served by nonstop service is considered, one of the most important factors is the expected market share which will be obtained by the carrier introducing the nonstop flight. Current market share models, which are designed for markets served primarily by nonstop service, include exclusively nonstop flights; the number of passengers flying on other flights is generally less than 20% and is considered negligible for most purposes. These market share models are therefore inapplicable to situations where the market share of carriers who do not offer nonstop service must be considered, and a model which takes into account other types of service is required.

4.1 Models

The most common market share model gives market share as a function of frequency share. This relation is given in Equation 4.1.

$$MS_i = \frac{Freq_i^\alpha}{\sum_j Freq_j^\alpha} \quad (4.1)$$

MS_i is the market share of Carrier i for a particular market; $Freq_i$ is the number of non-stop flights (or flights per time) of Carrier i in the market. α is an exponent representing the advantage of carriers with higher frequency shares. This nonlinearity is sometimes ex-

plained by the fact that most travelers plan round-trip itineraries, and are likely to fly on the same carrier in both directions. Because of this round-trip planning, carriers' market share depends to some extent on the number of combinations of departing and returning flight combinations they offer (the square of the frequency). Ultimately, the exponent is present in the model because the data justifies it; it typically has values between 1 and 2.

This model can be extended to include through and connecting flights by modifying the meaning of $Freq_i$ to include not only nonstop flights, but a weighted sum of all available flights. The model used here includes a term for flights with one stop and flights with one connection; flights with more than one stop, more than one connection, or both a stop and a connection are not included in the model. This is shown in Equation 4.2. The parameter b is the value of one-stop flights, and the parameter c is the value of one-connection flights. These are fractions of the value of a non-stop flight, and should in all cases be between 0 and 1.

$$Freq_i = Freq_{i_{nonstop}} + b \cdot Freq_{i_{onestop}} + c \cdot Freq_{i_{oneconnect}} \quad (4.2)$$

4.2 Methodology

The model given by Eq. 4.1 and 4.2 has been used historically, with coefficients $b = 0.4$ and $c = 0.1$. However, these values have not been verified in decades, and their applicability to the current market is questionable. It was therefore desirable to determine new values for these coefficients using recent market data.

An overview of the calculations is presented here.

4.3 Data Requirements

Two sources of data were required, one to provide information on the flights available and one to provide information on the passengers who flew on each carrier. For the flight data, the required information for each flight is:

1. Departure city

2. Departure time
3. Arrival city
4. Arrival time
5. Carrier
6. Flight number

The Airline Service Quality Performance (ASQP) data provided the required information. Other data sources, such as the Official Airline Guide (OAG) electronic data could also have been used; the ASQP data was selected as a matter of convenience.

For the passenger data, the required information is the number of passengers carried, broken down by OD market, carrier, and time. This information was provided by the O&D Plus 10% ticket sample database, which provides information on 1 out of 10 tickets issued, randomly selected. Both the O&D Plus and ASQP databases provided only information for domestic flights by the 10 Major carriers.

4.4 Market Selection

The first operation was the selection of markets suitable for inclusion in the calculations. Because the principal use for the model is to predict the effects of adding or removing nonstop service, those markets in which nonstop service was added or removed were selected. The markets selected are those which gained or lost nonstop service between 1995 and 1999, the years for which data was available. The markets must also have had at least one full year both with and without nonstop service during the period for which data was available, to reduce skewing of the results due to seasonal variations. The markets selected were also required to have had a total of 40 nonstop flights (both directions) during the period they were considered to have nonstop service; markets with a very small number of nonstop flights (such as one per week) were therefore excluded. These criteria served to exclude a number of markets which were served nonstop seasonally (winter or summer

only), and those which were served for only a few months before being abandoned. 123 markets which met the criteria were found.

4.5 Frequency Calculation

The number of nonstop flights was trivially determined; the number of flight legs which left from the one city of the pair and arrived at the other was counted. Determining the number of one-stop and connecting flights required examining all pairs of flight legs which lead from one city of the pair to the other, by way of any intermediate city (one carrier is examined at a time). If the flight number is the same, the flight legs are counted as a one-stop flight. If the flight numbers are not the same, the flight legs are examined for potential inclusion as a connecting flight. To be counted as a connecting flight, the second leg must depart at least 30 minutes, but not more than 120 minutes, after the arrival of the first leg. Using these fixed values for connecting times greatly simplified the problem compared to entering the actual connection times for each carrier and each airport, and is unlikely to significantly skew the results.

Frequencies were calculated by quarter; the limiting factor was the 10% ticket sample database, which provided data with a quarterly granularity.

4.6 Coefficient Calculation

Two methods were used for determining the coefficients b and c . A linear regression provided detailed error information with α constrained to 1; a nonlinear solution allowed α to vary but did not provide coefficient-specific error information.

4.6.1 Linear Solution

Linear regression “solves,” in a minimum-error sense, Equation 4.3.

$$Y = X \cdot c \tag{4.3}$$

c is an $N \times 1$ matrix of coefficients; it is the desired result. \mathbf{Y} is an $M \times 1$ matrix, and \mathbf{X} is an $M \times N$ matrix. When $M = N$, this equation has an exact solution; when $M > N$, linear regression is used to find the c which minimizes the error.

To perform a linear regression, the \mathbf{Y} and \mathbf{X} matrices must be generated. In this case, \mathbf{X} will be an $M \times 2$ matrix, as there are two coefficients in c . Rows in these matrices are generated using Equation 4.4, where one row is generated for each carrier i in each market.

$$\begin{aligned}
 & - \left(\sum_j MS_i \cdot Freq_{i_{nonstop}} - Freq_{i_{nonstop}} \right) = \\
 & b \cdot \left(\sum_j MS_i \cdot Freq_{i_{onestop}} - Freq_{i_{onestop}} \right) + \\
 & c \cdot \left(\sum_j MS_i \cdot Freq_{i_{oneconnect}} - Freq_{i_{oneconnect}} \right)
 \end{aligned} \tag{4.4}$$

Equation 4.4 results from straightforward algebraic manipulation of Equation 4.1 when α is 1. Rows which were entirely zero (representing a carrier which has no flights and no passengers in a particular market) were not included in the generated matrices.

4.6.2 Nonlinear Solution

The advantage of the nonlinear solution is that it removes the requirement that α be exactly 1. The disadvantage is that only an overall error is provided for the solution, and not a per-parameter error as is provided by a linear regression.

General nonlinear optimization algorithms generally take the form of maximizing or minimizing an objective function. In this case the objective is a function of α , b , and c . For each carrier in each market, the difference between the calculated market share from the given α , b , and c values and the actual historical value was found; the root mean square of this error value is the objective function to be minimized.

The specific algorithm used was the Nelder-Mead algorithm. This algorithm, which is general-purpose but not particularly efficient, was chosen for simplicity. Because this problem was a relatively tractable one, the slow running speed was not important. The

Table 4.1: Connecting Flight Value Linear Fit Results

	Value	Standard error	t-statistic
b (one-stop)	0.1748	0.0063	27.7
c (connection)	0.0129	0.0003	43.0

specific algorithm used should not affect the final result.

4.7 Results

4.7.1 Linear Results

The linear results, with an assumed $\alpha=1$, are presented in Table 4.1. The principal purpose of the nonlinear model is to examine the significance of the coefficients, and it can be seen that the fit is quite good and both coefficients are significant.

The value of connecting flights is much lower than that of non-stop flights, even though it is not clear that these flights take longer to complete. Part of this is undoubtedly due to the fact that infeasible connections - those which go across the country and then back - are included, even though it is unlikely that a significant number of passengers will fly them. The presence of one-stop flights before connecting flights on the CRS menus undoubtedly explains part of this as well.

4.7.2 Nonlinear Results

The nonlinear results are presented in Table 4.2. The nonlinear solution does not provide per-parameter errors statistics like a linear regression does; the total root mean square error, calculated as described in section 4.6.2, was 0.144.

The nonlinear results provide higher values for the b and c coefficients than the linear solution. This is likely a consequence of the restrictions on α present in the linear solution. Because carriers with nonstop flights also tend to be the ones with high market shares,

Table 4.2: Connecting Flight Value Nonlinear Fit Results

Parameter	Value
α	1.107
b	0.233
c	0.048

the nonlinearity which is an effect of α was artificially captured as low values of b and c , which also served to further increase the market share of those carriers with above-average marketshares.

Chapter 5

The MIT Extensible Air Network Simulation

The MIT Extensible Air Network Simulation (MEANS) is an event-based simulation tool which can be used to analyze current and hypothetical network configurations and rules.

The discussion of MEANS is divided into three sections: data sources, calculation methods, and results.

5.1 Data Sources

Several types of data are required by the simulation. These include a schedule of flights, airborne rates and times, arrival and departure rates at airports, and time required for taxi and other ground events. Each of these is addressed below.

5.1.1 Schedule of Flights

The schedule of flights for historical days is obtained from the Airline Service Quality Performance (ASQP) database. This database was selected over sources such as the Official Airline Guide (OAG) data because it has the highly desirable property of including the registration number of the aircraft. The aircraft registration number is crucial for following an aircraft through its flights and correctly modeling the propagation of delays.

The major limitation of the ASQP database is that it includes information only for domestic jet flights operated by the ten US Major airlines. The OAG database includes all scheduled commercial flights, but still does not include freight or general aviation traffic. To address these data limitations, a count of all operations in 15-minute periods was obtained from the FAA's Consolidated Operations and Delay Analysis System database[2]. This count enabled flights to be generated and added to the flights in the ASQP database, bringing the total level up to the historical value.

5.1.2 Airborne Performance

In the US, the principal cause of delays is a lack of capacity at airports, not in the air. MEANS therefore uses a simple model based on historical information. The average airborne times between two cities is used when historical information is available. Where historical information is unavailable, an approximate airborne time is calculated from distance and scheduled time using a formula developed as a "best fit" from those cases where actual data was available.

It is also possible to vary airborne times around the average, but the results of the simulation are generally quite insensitive to small changes in airborne time. It might be possible to obtain better results by segregating the results by aircraft type, but this approach was not pursued for reasons of tractability.

5.1.3 Ground Operations

Times required for ground operations are determined in a manner similar to those used for airborne times, described above. Times are calculated by airport, but not by aircraft type or configuration for reasons of data availability and numerical complexity. [11]

5.1.4 Weather

Historical weather information from the CODAS database is used as weather input. Historical data has the advantage that it captures the correlations between proximate cities; alternative approaches such as a Markov-chain generated weather model do not capture

this. The CODAS data provides wind and visibility data from which it is possible to determine the operating condition of the airport, as well as a note as to whether the conditions were Instrument or Visual Meteorological Conditions.

5.1.5 Airport Arrival and Departure Rates

The calculation of the arrival and departure rates has been one of the most difficult and involved parts of the MEANS development. As a result of this, several different sources of data have been used. Historical data from tower logs at the Boston Logan airport were used in the initial development of the simulation, as these provided the most detailed and accurate information. These were later replaced with capacity envelopes from the FAA Benchmark report. [1]

5.2 Calculation Methods

The central data structure in MEANS is the flight leg. Flight legs are updated and modified as flights are rescheduled or delayed, and are passed between different parts of the program as the flight moves. There are a number of states in which a flight leg can be. It can be inactive (another flight leg with the same aircraft is still in progress), at the gate, taxiing out, waiting to depart, in the air, waiting to arrive, and taxiing in. Timing is controlled through a queue of events; any of the modules can schedule an event to occur later. Having a routine which handles one event and schedules another for later is the most common way of having recurring actions.

5.2.1 Flight Data Structure

For each flight leg, an airline, flight number, tail number, departing and arriving city, and scheduled, rescheduled, and controlled departure and arrival times are stored. Additionally, the actual time of all state transitions is recorded as the flight leg changes states. Time is recorded as seconds after a defined point. For a single-day situation, this can be seconds after midnight; for multiple day runs, seconds since the Unix epoch (midnight January 1,

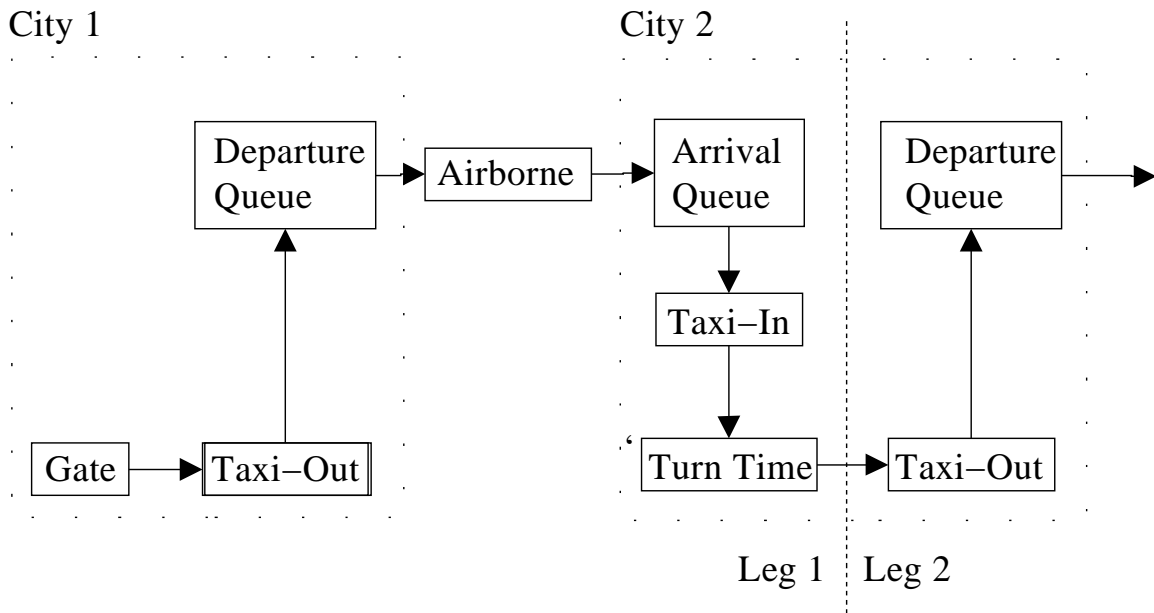


Figure 5-1: MEANS Flight Progression

1970, GMT) is a convenient choice. Figure 5-1 shows the states through which a flight passes.

5.2.2 Schedule Loading and Generation

Because some flights do not have full information in the ASQP data and must be added based only on operation counts, a special case for the format must exist to mark these flights. Special symbols for airline and departure/arrival airports are used to identify these “padding” flights, and allow them to be processed properly.

The schedule is read from a custom-format data file, which includes exactly the information required, and written to another simply formatted file which includes the originally scheduled information plus the information about actual times. This is the principal output of the simulation, and most results are obtained by postprocessing this file.

For the flight legs without complete information, however, a full schedule file is not available to the simulation. Instead, an abbreviated description is used, which specifies the number of arrivals and departures to be added in a given time period at a given city. Several methods are available for specifying this information. The additions can be specified as a

total number of legs to add, a total target number (in which case the number to add will be calculated based on the number of flights already present in the schedule), or a percentage of the scheduled flights which should be added.

One piece of information which cannot be left out, even for these flights, is the aircraft tail number. This is vitally important to obtain proper delay propagation. Furthermore, assigning each leg a unique tail number is an inadequate approach. In a ground delay program, only arrivals are controlled. Since arrivals become departures, the number of departures must be reduced when the arrivals are reduced. However, this can not be achieved if some of the departures are not recognized as being the same aircraft as the arrivals. Thus, the algorithm which inserts “fake” flights attempts to match incoming and outgoing tail numbers, based on assumptions on the minimum turn time of the aircraft.

5.2.3 Taxi and Airborne

Historical information for taxi and airborne times is read from a datafile. The data file presents information by percentile, and an appropriate percentile or distribution is selected by the simulation. Taxi data will always be available, but airborne data may need to be calculated. In this case, information about the distance between the two cities and the scheduled duration are used to calculate an approximate airborne time.

When an aircraft is passed to the taxi or airborne handler, the handler calculates the time required and schedules an event at the appropriate time in the future to pass the aircraft to the next state.

5.2.4 Tower Controller

The tower controller, which passes aircraft from the departure queue into the air and from the arrival queue onto the ground, is one of the most complex parts of the simulation. It is responsible for metering the flow of aircraft and ensuring that appropriate separation standards are not violated.

In its most basic implementation, the tower is supplied with a fixed arrival and departure rate, from a source such as historical data. It schedules arrival and departure events to occur

at times evenly spaced throughout the hour, and whenever there is a slot for an aircraft to depart or arrive, the next aircraft in the queue is processed. No trading between departures and arrival is possible.

A more sophisticated implementation of the tower uses a Pareto frontier showing the number of possible arrivals for a given number of departures. In this case, the tower controller must periodically select an operating point on the curve (an action notionally corresponding to selecting a configuration or runway allocation). The algorithm used to calculate this operating point has significant effects on the behavior of the system.

Presently, the operating point is chosen using information about both scheduled flights in the future and flights currently waiting. The number of departures and arrivals expected in the next period (often one hour) are calculated, and added to the number of departures and arrivals currently waiting in the queue. An operating point is selected such that the ratio of the actual departures to arrivals is the same as the ratio of desired departures to arrivals. This has worked well in practice.

A given airport will have data for multiple Pareto frontiers, corresponding to operating states such as Visual or Instrument Flight Rules (corresponding to the Visual and Instrument Meteorological Condition weather conditions). Other specialized local configurations may also sometimes occur, such as differing capacities as a function of wind direction. Somehow the simulation must select one of these to be in effect at any given time. The simplest way of doing this again requires the configuration of all airports to be directly specified in a data file; this is most useful for modelling exact historical situations where such information is available. An alternative approach, which is capable of selecting between VFR and IFR configurations for each airport, calculates the conditions based on visibility and ceiling information from the weather at the airport in question.

5.2.5 The FAA and Ground Delay Programs

When delays become particularly bad at an airport, a Ground Delay Program is put in place to hold aircraft travelling to that airport on the ground, as waiting on the ground is safer and less expensive than waiting in the air. Thus, MEANS has a ground delay program

implementation which models the FAA in instituting ground delay programs.

In its simplest form, a datafile can be provided with historical ground delay program information. This specifies the beginning and end of the ground delay program, the arrival rates, and the time at which the program is announced.

A more sophisticated approach uses predicted airport capacities (based on predicted weather) to automatically institute a GDP when demand exceeds capacity by a certain amount.

Once a GDP has been initiated through any of the above methods, it must be performed. This involves generating and assigning slots to flights. The algorithm used by MEANS is a simplified implementation of the ration-by-schedule algorithm with compression, outlined below.

First, a list of flights is generated with originally scheduled arrival times in the window over which the GDP is operating is generated; this list is sorted by arrival time. A table of slots is then generated based on the predicted arrival rate. The slots are allocated to airlines in the same order as the originally-scheduled flights, up to the number of slots in the table (there will obviously be fewer slots than original flights). This is the ration-by-schedule algorithm.

Each slot then has a flight assigned to it. This is done by finding the next available flight from the airline to whom the slot belongs and assigning it to the slot. This examination of the schedule of flights includes any delays the airline has already made to the schedule, unlike the original slot assignment. Thus, it is possible that it will not be possible to fill a slot with a flight from the desired airline, if the only candidate flight has been delayed until after the time of the slot. In this case, a flight from another airline is shifted forward, with the displaced airline getting priority for the slot opened up by the move. This is the compression algorithm.

Once the initial assignments have been made, the airline agents are notified, and allowed to swap and cancel flights. When a flight is cancelled and the airline cannot find another flight to assign to its slot, the compression algorithm is rerun.

5.2.6 The Airline

In practice, an airline will decide which flights should suffer the most based on criteria such as loads, missed connections, or maintenance requirements. However, data for these considerations is not available to MEANS. At present, therefore, the airline agents simply cancel flights which will be delayed more than a specified maximum amount by the GDP. Two hours has been used as this maximum for work done to date.

5.3 Results

In testing, MEANS was evaluated in two scenarios - one in which heavy delays at one airport made the rest of the system irrelevant (for which Boston on January 28, 1999 was an example), and a summer day when the overall load was high but no particular locations had a dominant delay (an example of which was July 30, 1999). The simulation performed well in the former case and reasonably in the latter.

Figure 5-2 shows the aircraft in the arrival and departure queues at Boston Logan airport on January 28, 1999. Boston had severe inclement weather on this day, and a ground delay program was in effect most of the day. As described in section 5.1.1, the flights with information obtained from the ASQP database were supplemented with generated flights from the ETMS data. In order to allow for a valid comparison with the historical ASQP queue sizes, these generated flights were not included in the count of the aircraft in the queue. As can be seen from the figure, the MEANS results track the historical results from the ASQP data quite well.

Figure 5-3 shows the simulated and actual arrival times of flights arriving at Boston on this day.

Figure 5-4 shows results for a few sample cities for the case of July 30, 1999, when there was good weather and generally heavy traffic across the system. The results are better for some cities than others. Particularly, those cities with more straightforward operations and fewer operating conditions are better simulated.





Figure 5-2: MEANS Queue Sizes, Boston, January 28, 1999

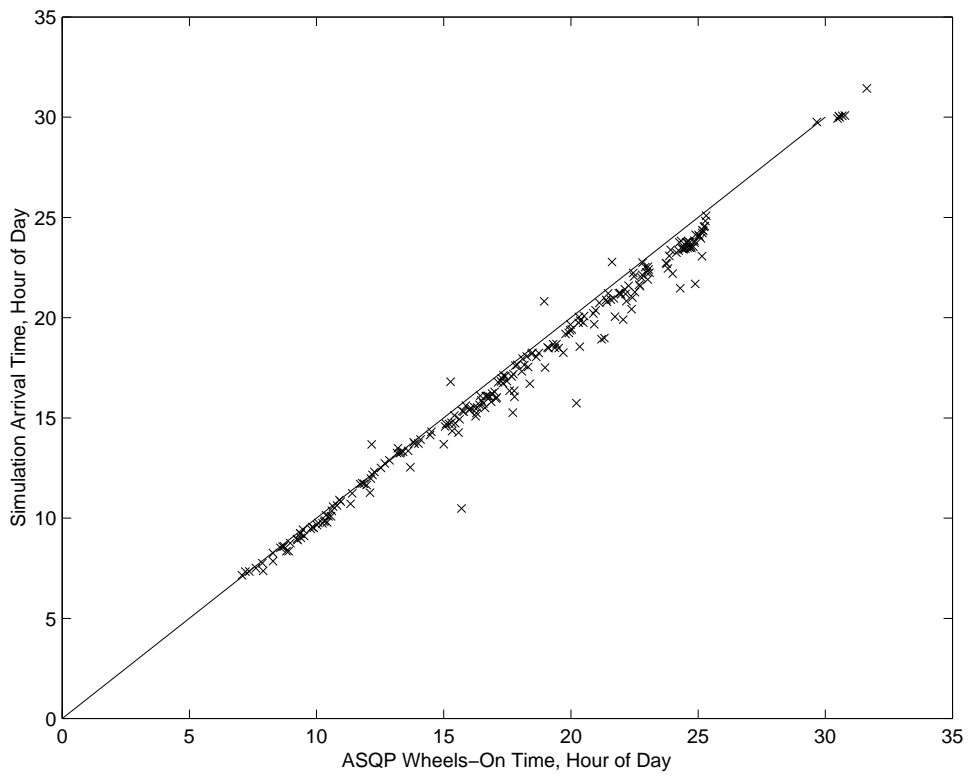


Figure 5-3: MEANS Arrival Times, Boston, January 28, 1999

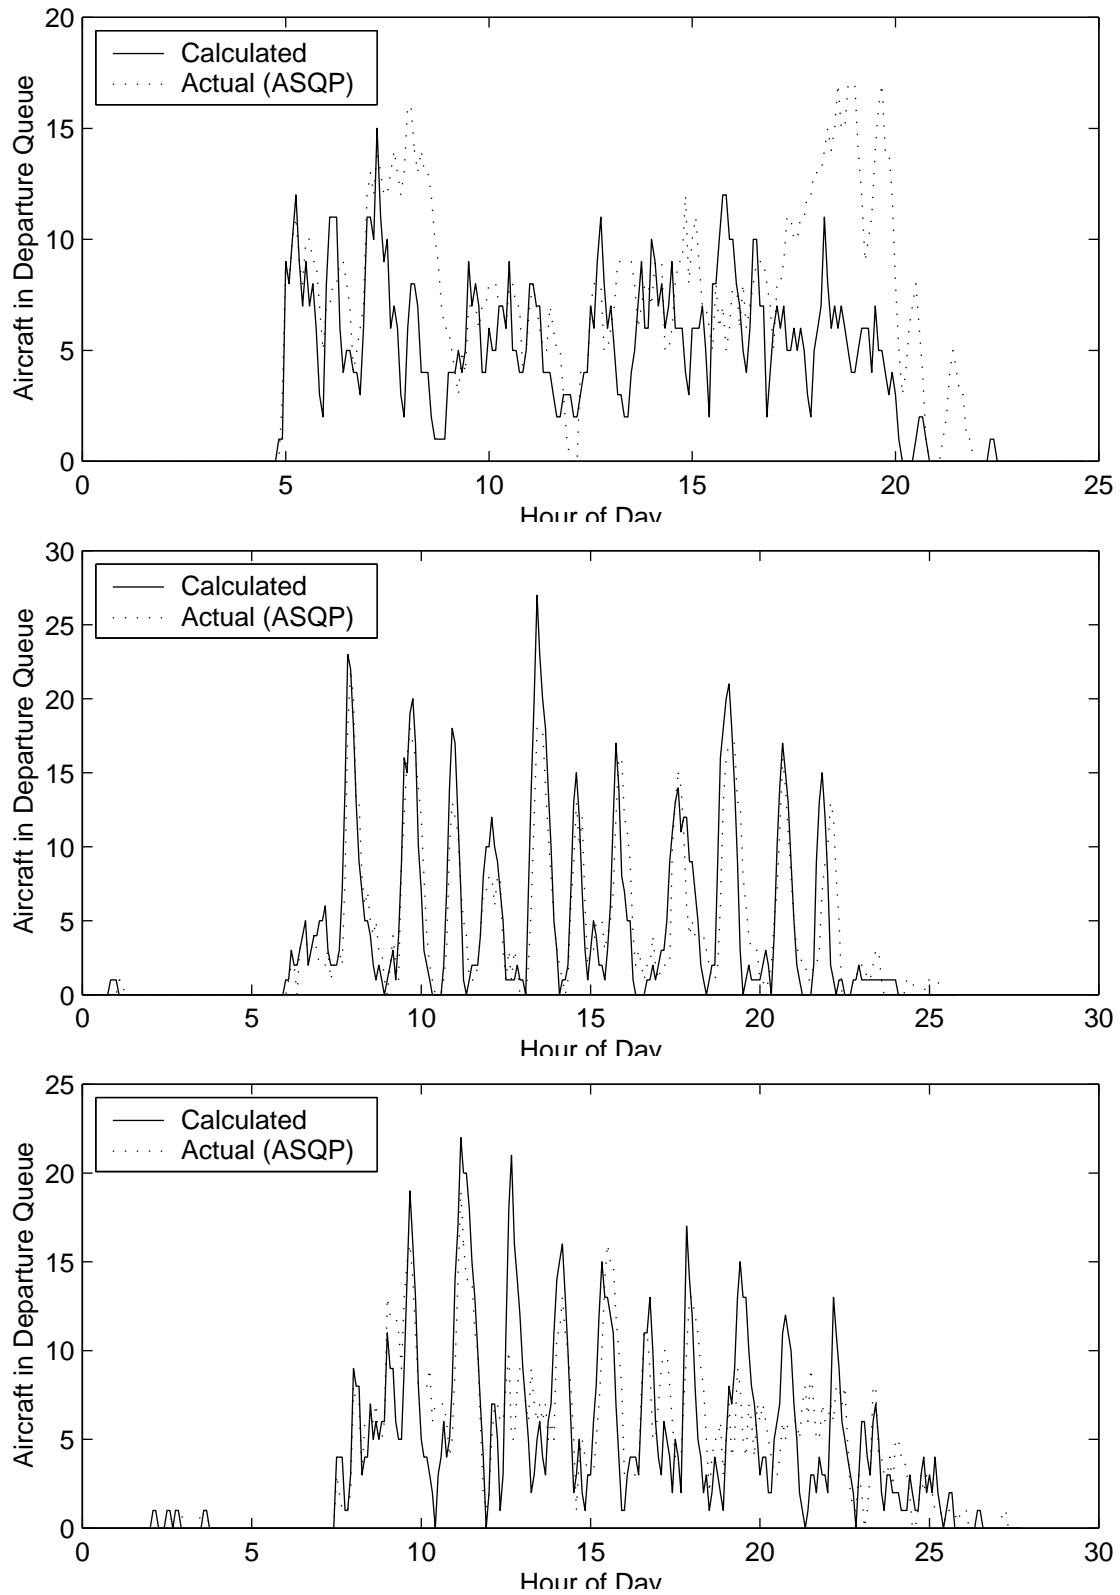


Figure 5-4: MEANS Results, BOS, IAH, and PHX, July 30, 1999

Chapter 6

Approach

It is not possible to calculate the effect of all possible schedule alterations in full detail due to the computational complexity of such a task. Therefore, a two-stage approach is used, in which preliminary calculations of lower accuracy are performed to identify plausible flights for replacement, and then a full calculation is performed using this set of candidates. The preliminary calculations ignore the network-wide effects of changes, in order to avoid recalculating the entire scenario.

This calculation procedure is shown in Figure 6-1. Each step is described below.

6.1 Cost Calculations


6.1.1 Determination of Hourly Costs

The hourly cost for an aircraft type is calculated from Form 41 data, as described in section 3.2.1. Two values are calculated, for time spent on the ground and time spent in the air, using costs as shown in table 3.1. Because the purchase and sale of aircraft is not being considered, ownership costs are not required.

6.1.2 Determination of Aircraft Type

The type of aircraft is not actually included in the ASQP data. It must be looked up separately. The JP Fleet database was used to do this.

Preliminary Calculations



Final Calculations

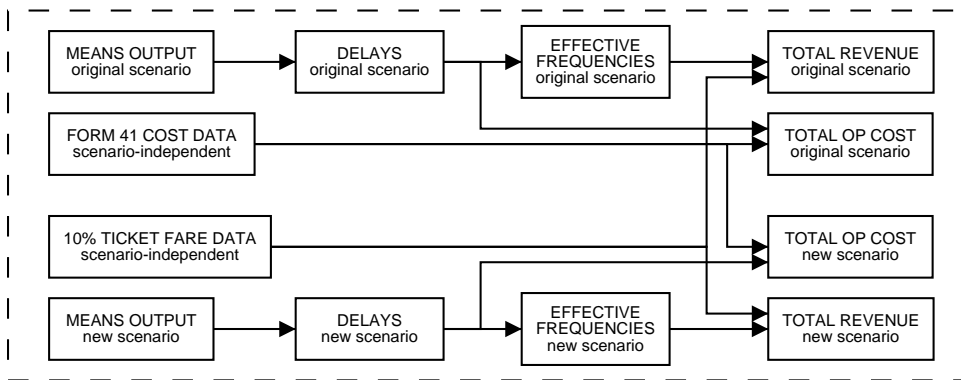


Figure 6-1: Information Flow

However, the ASQP data contains some errors in its reported tail numbers, and some aircraft are missing from the JP Fleet database. When the tail number of an aircraft could not be found, an attempt was made to determine the aircraft type by looking at aircraft which flew the same leg on other weeks in the same month. Of the 9.4% of the data which could not originally be looked up, over $\frac{2}{3}$ were amenable to this technique, leaving only 2.7% unknown. For these few, an average value for the carrier's known fleet was used.

6.1.3 Preliminary Cost Calculations

The cost of operating a flight, from the point of view of preliminary analysis, is taken to be only the direct operating costs of the flight, calculated from the amount of time it spends in various states. The indirect effect of changed costs in other flights due to delays which would be reduced or increased by changes to the flight in question were not taken into

account.

For the case of an existing flight which is being considered for removal, the operating cost of the flight is calculated from the simulated times it spends on the ground and in the air, including delays. This will cause heavily-delayed flights to be favored for removal, as desired.

For the case of a flight which has not been flown historically or simulated and is being considered for addition, no delay information is available, so the flight is analyzed based on the schedule. Because the focus is on adding flights which will not be severely delayed, the cost due to delay is small compared to the base cost calculated from the schedule time. Delays will be taken into account in the final analysis.

6.1.4 Final Cost Calculations

For the final cost calculations, the total cost of two complete scenarios are being compared. These costs are calculated by summing the costs of all flights for all operations, not just those flights which differ between the scenarios. This fully captures the indirect effects of delay reduction, wherein cancelling one flight causes other flights at the same airport to be delayed less due to reduced demand.

This calculation is more accurate than the preliminary calculation because it captures the network-wide effects. However, it also makes it difficult to analyze the effects of individual flights. Thus, while the final calculation produces better overall results, it is less useful in providing an intuitive understanding of the phenomena.

6.2 Revenue Calculations

6.2.1 Frequency Calculation

Frequencies are calculated using the fundamental approach described in Chapter 4 for comparing nonstop frequencies with those of one-stop and connecting flights. Connecting and one-stop flights are weighed as described in sections 4.6.2 and 4.7.1, which describe the nonlinear solution procedure. One additional consideration is the perceived devaluation of

flights based on delay. Previous work has shown that passengers perceive flights as becoming less valuable by 2.5 divided by the length of the flight in minutes for each minute of mean delay. [10] This is shown in Equation 6.1

$$value = \frac{NormalDuration}{NormalDuration + 2.5 \cdot Delay} \quad (6.1)$$

The effect of this is to cause passengers to favor flights with less delay over flights with more delay. The relative value of connecting flights versus nonstop flights is largely unchanged, as this effect is already captured by the general weighting of connecting and one-stop versus nonstop flights, described in section 4.6.2. The significance of this delay-based devaluation is that it causes heavily delayed connecting flights to be perceived as less valuable than less-delayed ones, an effect which is not captured by the general weighting applied to all connecting flights. This decrease in effective frequency reduces the revenue obtained from delayed flights in the base case, due to decreased passenger preference. This causes the revenue loss from cancelling connecting flights through extremely congested and heavily delayed hubs to decrease, because these flights were worth less in the base case and thus had less to lose.

6.2.2 Preliminary Revenue Calculations

Revenue is calculated by assigning market share using frequencies, as described in section 4.1. When a flight is considered for cancellation, the effect on the carrier's frequency in all markets served by that flight is calculated. The cost of cancelling the flight is taken to be the difference in the revenue obtained from the set of frequencies which includes the flight in question and the set of frequencies which does not.

When a flight is considered for addition, the expected revenue gain is calculated in a similar manner. However, the new flight is not considered for part of a connection; only the increased frequency in the nonstop market served directly is considered. Because the delays are not recalculated for hypothetical cancellations in the preliminary analysis, the delay-weighted frequencies cannot be updated.

For simplicity, each aircraft is assigned to exactly one city-pair, flying back and forth

throughout the day. Because most city-pairs selected do not have existing nonstop service, it is not possible to use historical times; theoretical times calculated from distance are therefore used. In addition to the airborne times calculated from the distance, average taxi times are added.

The city pairs considered here are too close for a “red eye” overnight flight to be of interest, so it is assumed that a 14-hour window is available in which to schedule flights. From the leg length calculated as described above, the number of legs which can be flown in one day is calculated; a universal 40-minute time on the ground between legs is assumed. This number of flights is used to generate new frequencies, which in turn produce a market share. This market share is combined with the total demand and average fare for the market to produce a revenue. A check is performed to ensure that the capacity of the aircraft at a reasonable load factor is not exceeded.

6.2.3 Final Revenue Calculations

In the final revenue calculations, a new set of weightings is calculated from the actual delays of running the new scenario. These are then used in the revenue calculations for the new scenario, fully capturing the network-wide effects of the delays. The expected market-share-based revenue is then calculated for all flights in the original and new scenarios, not just those which were modified.

This network-wide approach includes the effects of new delay weightings for flights. This means that additional revenue realized by reducing delays on those hub flights left in place as a consequence of reduced demand at the hub is included. So is any revenue from connections which may have been introduced as a side effect of the new nonstop flights.

Like the final cost calculations, the final revenue calculations provide a more accurate result, but also make it impossible to examine the effect of an isolated change.

6.3 Scenario Generation

In order to perform the final calculations, MEANS must be run on a hypothetical scenario generated from the preliminary calculations. The generation of this scenario consists of

removing those flights which are being replaced, and inserting those which are being added. Flight durations are calculated as described in section 6.2.2. Flights are started in an east to west direction due to time zone considerations, and simply fly back and forth between the designated cities with the specified number of flights.

MEANS is rerun with this replacement schedule, and all of the results which depend on the MEANS output are recalculated, to allow the full network-wide effects to be included in the final calculations.

Chapter 7

Results

Continental's flights in the Newark and Houston hubs were considered for replacement using the methodology described in Chapter 6. Because many flights visit both hubs in the same day, it was not possible to isolate one hub. Results are therefore presented for both hubs.

7.1 Preliminary Results

In the preliminary calculations, a set of 59 flights to be replaced were identified, with a more profitable replacement for each. Table 7.1 shows the average effect of these replacements. All values are in dollars, calculated on a daily basis. For comparison, results which include the effects of delays and results which do not include the effects of delays are presented, to show how much of the improvement is possible only because of delays in the current situation.

The revenue loss is the amount of revenue which is lost from the cancellation of the existing legs; the revenue gain is the revenue obtained from the addition of the new legs. The case which does not include the effect of delays ignores the perceived decrease in frequency due to the devaluation of delayed flights; the case which considers delays includes this effect.

The operating cost change is the difference in cost between the old and new scenarios. In the case which does not include the effect of delays, the operating cost is calculated based

Table 7.1: Average Results of Preliminary Calculation for Continental

Without effect of delays	
Revenue Loss	13015
Revenue Gain	20726
Op Cost Increase	2206
Net Profit Increase	5504
Including effect of delays	
Revenue Loss	12128
Revenue Gain	23004
Op Cost Increase	1886
Net Profit Increase	8990

on the scheduled times; the case which considers the effect of delays includes the extra cost arising from delays. The positive number indicates a higher cost in the new flights than the old. This is due to the increased utilization which is obtained as a consequence of no longer needing to time flights to meet banks of connections. Complete data is given in Appendix A.

Some of the flights which are suggested for replacement would have been desirable (although less so) to replace even without delays, but many flights become desirable to replace only due to delay effects. Approximately 1/3 of the replacements would have had a detrimental effect were it not for the increased costs and decreased revenues caused by delays in the current situation.

The total expected benefit from the preliminary calculations, from both reduced costs and increased revenues, is approximately \$530,000. When we examine the original situation without delays, the situation is better and allows for less improvement. In this case, a profit improvement of \$325,000 could have been obtained. The difference between these, over \$200,000 per day, represents a delay-induced loss in the current situation which can be recovered by increasing nonstop flights.

Table 7.2: Comparison of Delays in Original and Improved Case for Continental

	Original Case	Replacement Case
EWR Operations	450	399
EWR Total Outgoing Delay, Minutes	1030	832
EWR Total Incoming Delay, Minutes	651	530
EWR Average Outgoing Delay	4.6	4.2
EWR Average Incoming Delay	2.9	2.6
EWR Flights Delayed 15 Min	17	18
IAH Operations	595	502
IAH Total Outgoing Delay, Minutes	1553	1067
IAH Total Incoming Delay, Minutes	754	448
IAH Average Outgoing Delay	5.2	4.4
IAH Average Incoming Delay	2.5	1.7
IAH Flights Delayed 15 Min	35	26
Systemwide Average Outgoing Delay	4.9	4.5
Systemwide Average Incoming Delay	2.0	2.0

7.2 Final Results

A hypothetical schedule was generated based on the flight replacements suggested by the preliminary calculations, and a full calculation based on a MEANS run using the new schedule was performed.

The total number of flights at both the Houston and Newark hubs significantly reduced, and so was the delay. Table 7.2 shows several statistics comparing the original and new scenarios.

The overall results, in the form of operating cost and revenue for Continental, are shown in Table 7.3. The net improvement was \$551,000, which is very close to the preliminary estimate of \$530,000.

7.3 Conclusions

It appears that there are significant monetary gains to be obtained from cutting back on hub flights in favor of nonstops, and that additional costs and lost revenue from delays contribute significantly to this effect. As hub congestion increases, delays will increase,

Table 7.3: Final Calculation Results

	Historical Scenario	Proposed Scenario	Change
Operating Cost	3538102	3766285	228183
Revenue	6373835	7152746	778911
Operating Contribution	2835733	3386461	550728

and the motivation to add nonstops will also increase.

A savings of \$550,000 per day translates to over \$49 million per quarter. For comparison, Continental reported a net income of \$42 million for the second quarter of 2001.[4] Thus, while it unlikely that the entirety of the savings predicted will be realizable due to additional constraints beyond the scope of the model, the potential improvement from increased nonstop routings clearly merits serious consideration.

Chapter 8

Directions for Future Work

There are several directions in which the work done here could be extended. These relate primarily to additional capabilities in the MEAN simulation and additional calculations relating to the replaced flights directly.

8.1 Other Airlines

A sample case for Continental was calculated here. A similar approach could be used to calculate improvements for several of the other major carriers with a significant hub and spoke network. This would allow for a comparison of hubs amongst carriers.

8.2 MEANS additions

There are several planned additions to MEANS which would allow additional work to be done. In addition to general improvements in accuracy, MEANS is looking to integrate with outside decision-support tools which will improve the decisions made by the hypothetical airlines in the simulation. Additionally, the capability to route aircraft intelligently for prolonged periods, including maintenance and disruption recovery, should allow the simulation to be run for an entire season or entire year instead of a small set of representative days.

8.3 More Sophisticated Scenarios

The replacement scenario generated here required that each new flight go back and forth between two cities. A more sophisticated approach with greater flexibility might provide superior results. This is particularly the case if a longer run of MEANS including maintenance and crew scheduling requirements is also included.

8.4 Enhanced Calculations

The present calculations identified a set of flights in the preliminary stage and considered a scenario replacing these flights in the final stage. A more sophisticated calculation might be able to generate a scenario for all possible combinations of flights. If a few of the flights were actually unprofitable in the final analysis in spite of being profitable in the preliminary one, this approach would eliminate them, providing an optimum answer.

8.5 Shared Hubs

Some airports are hubs to multiple large airlines. Chicago O'Hare, for example, is a significant hub for both American Airlines and United Airlines. At these airports, the situation is more complex, because if one carrier removes traffic, the other carrier may add additional flights to fill up the space. This could ultimately lead to the carrier attempting to reduce its hub operations being shut out of the hub entirely. A game-theoretic analysis applied to this case could answer the question of whether it would be desirable to reduce flights at a shared hub, or whether only hubs with one dominant airline benefit from such a reduction.

Appendix A

Formulae, Values, and Data

A.1 Fits to Historical Data

A.1.1 Actual Flight Time

Two different formulae for flight time are available, one for the case where a “scheduled time” is available (used in MEANS to calculate flight time for city pairs without historical data), and one for the case where it is not. Equation A.1 shows the calculation of actual airborne time from published gate to gate times; values are in minutes. Equation A.2 shows the calculation from distance when a scheduled time is not available. Distance is in miles; times are in minutes.

$$ActualAirborne = Scheduled \cdot 0.934 + 16.3 \quad (A.1)$$

$$ActualAirborne = Distance \cdot 0.1129 + 29.6 \quad (A.2)$$

Figure A.1.1 shows the actual airborne times and those calculated from distance for cities where both are known. These cities are the dataset with which the fit was performed.




Figure A-1: Airborne Time Calculated from Distance

A.2 Detailed Continental Results

Detailed results from the flight replacements for Continental, described in section 7.1, are shown below. Table A.1 shows the flights which were replaced and their replacements. Table A.2 shows the per-flight results of the preliminary calculations using these flights. Table A.3 is a list of city abbreviations used in the other tables.

Table A.1: Replaced Flights for Continental

Flight	Replacement Cities	Original Cities
N76073	LAX-MDW	HNL,IAH
N19072	MCO-MDW	HNL,IAH
N29124	MCO-SJU	IAH,LAX,SFO
N14121	DEN-MDW	EWR,IAH,SEA
N33132	MCO-MKE	EWR,IAH
N17122	MDW-SFO	ANC,EWR,IAH,LAX,SEA
N17104	LGA-MKE	EWR,SEA
N68047	JFK-STT	EWR,IAH,LAX
N78005	BOS-RDU	EWR,IAH
N18112	FLL-MDW	EWR,PBI
N17326	LGA-RDU	BOS,IAH
N29717	BOS-MKE	EWR,SNA
N16884	DFW-MKE	EWR,IAH,SFO
N13891	EWR-MKE	IAH,MIA,TPA
N21108	DFW-MDW	EWR,FLL,IAH,LAX
N16703	LAX-MKE	EWR,SNA
N14381	LIH-SFO	DCA,EWR,IAH
N17233	MKE-PHL	EWR,SAN
N69602	MDW-PHX	BNA,EWR,IAH
N14115	DCA-MKE	EWR,LAX
N14818	MDW-RSW	IAH,MIA,PHX
N24706	DTW-RSW	EWR,IAH,PDX,SJU
N12225	LGA-MDW	CLE,EWR,IAH
N17663	DFW-LGB	BNA,CHS,EWR,IAH
N14653	DEN-MKE	BWI,CLE,LAX
N13720	SEA-SNA	ABQ,IAH
N33817	CVG-PHL	DCA,EWR,IAH
N14102	IND-RSW	EWR,LAX
N10801	FLL-JFK	DTW,EWR,PBI
N511PE	BUF-LGA	ATL,BWI,IAH
N12327	CVG-EWR	IAH,MSY,SEA
N72825	LGA-STL	BOS,DTW,EWR,IAH
N70353	CVG-LGA	ATL,IAH,TUS
N14219	BWI-SAN	EWR,IAH,SJU
N13227	KOA-SFO	EWR,IAH,SAN
N38727	LAX-MRY	EWR,IAH
N14346	LGA-MCI	DFW,EWR,IAH
N16642	MKE-SFO	DCA,IAH,RIC
N14106	CMH-LGA	IAH,LAX,MCO,MSY
N57837	HOU-LAX	IAH,MIA,MSY,SAN
N14668	IND-LAS	ELP,IAH,RDU
N26210	FAT-LAX	EWR,IAH,LAX
N72830	DTW-FLL	DTW,EWR,TPA
N12811	LAX-OGG	EWR,IAH,TPA
N14601	DCA-HSV	AUS,EWR
N24715	BWI-SFO	DCA,IAH,SAT
N33608	PSP-SEA	ATL,COS,IAH
N33785	CLE-MKE	IAH,LAS,ORD
N16217	LAS-MDW	BUF,EWR,SEA
N16806	DTW-TPA	DEN,EWR,MIA
N47332	JFK-MDW	BDL,CLE,IAH,LAS,MFE
N69803	ORD-RSW	EWR,IAH,TPA
N12218	FLL-ORD	IAH,LGA,TPA
N69351	MIA-TPA	EWR,IAH,IND,RSW
N26232	MIA-TLH	BOS,CLE,EWR,IAH,SEA
N17644	BWI-OAK	BNA,EWR,MSP
N14662	OGG-SFO	IAH,RNO,SNA
N27722	CMH-MDW	CLE,EWR,IAH,SFO,SJU
N19382	JAX-LGA	IAH,MSY

Table A.2: Preliminary Calculation Results by Flight

Existing Flight	New Cities	Without Delay Effects				With Delay Effects				Difference due to Delays
		Op. Cost Change	Rev. Loss	Rev. Gain	Total Change	Op. Cost Change	Rev. Loss	Rev. Gain	Total Change	
N76073	LAXMDW	8952	3423	38861	26485	8797	3121	43973	32055	5414
N19072	MCOMDW	3539	3423	39385	32423	3523	3434	43232	36274	3835
N29124	MCOSJU	1717	16444	40413	22251	760	12023	42197	29414	6205
N14121	DENMDW	-1380	9869	30950	22461	-1562	9912	32112	23762	1118
N33132	MCOMKE	11839	1033	28313	15440	11794	957	31734	18982	3497
N17122	MDWSFO	-8294	28306	28465	8453	-9225	26975	30688	12937	3553
N17104	LGAMKE	-3995	15191	26384	15188	-4240	16042	29685	17882	2449
N68047	JFKSTT	-5141	33482	28934	593	-5850	30481	29510	4879	3576
N78005	BOSRDU	4242	8329	25273	12701	4114	8308	29489	17067	4237
N18112	FLLMDW	15006	1844	28134	11284	14980	1560	29086	12546	1235
N17326	LGARDU	412	6773	22747	15562	283	6308	26672	20080	4389
N29717	BOSMKE	4605	5755	24275	13914	4582	5642	26373	16148	2211
N16884	DFWMKE	-639	14103	23713	10249	-1115	12288	26343	15170	4445
N13891	EWRMKE	5129	7192	24125	11803	4685	5047	26323	16590	4343
N21108	DFWMDW	-3220	27120	23923	23	-4707	21845	25400	8262	6752
N16703	LAXMKE	823	12896	20087	6368	684	12227	23956	11045	4537
N14381	LIHSFO	7742	6330	23347	9273	7590	4914	23896	11391	1965
N17233	MKEPHL	2822	6991	21627	11813	2768	7224	23753	13760	1892
N69602	MDWPHX	6549	6174	20994	8269	6310	5274	23610	12024	3515
N14115	DCAMKE	3221	10841	21005	6943	3196	12129	23316	7990	1022
N14818	MDWRSW	3508	14211	22270	4550	2762	10361	22855	9731	4434
N24706	DTWRSW	-3336	15790	20266	7812	-3434	15810	22556	10180	2269
N12225	LGAMDW	3089	8485	20359	8784	3016	8322	22056	10717	1859
N17663	DFWLGB	5431	9361	20672	5879	5335	8898	21848	7614	1639
N14653	DENMKE	-2068	15513	18409	4964	-2399	13832	20691	9257	3962
N13720	SEASNA	5777	6198	17084	5107	5733	6496	20670	8440	3288
N33817	CVGPHL	4568	6999	15850	4283	4401	7025	20559	9132	4682
N14102	INDRSW	-4214	24631	18917	-1500	-4704	23617	20417	1503	2513
N10801	FLLJFK	3845	12858	18136	1432	3316	10842	20246	6087	4126
N511PE	BUFLGA	587	13880	17050	2583	273	13179	20225	6773	3876
N12327	CVGEWR	2373	10199	16224	3651	2012	10167	20121	7942	3929
N72825	LGASTL	636	14868	15031	-473	-41	12953	19903	6991	6786
N70353	CVGLGA	3077	10271	15987	2637	2809	9901	19795	7085	4179
N14219	BWISAN	3476	10804	15411	1130	3410	10817	19774	5546	4349
N13227	KOASFO	-1185	16225	17473	2433	-1387	16009	19750	5128	2492
N38727	LAXMRY	-312	11432	18870	7749	-424	10606	19743	9561	1699
N14346	LGAMCI	6619	7875	16474	1979	6484	7996	19697	5215	3101
N16642	MKESFO	4238	13326	17260	-303	4038	13175	19676	2462	2566
N14106	CMHLLGA	-4150	22675	16070	-2454	-5109	22339	19399	2169	3664
N57837	HOULAX	-51	20145	16656	-3436	-1006	16131	19378	4253	6735
N14668	INDLAS	7117	8446	17187	1623	7029	8736	19364	3598	1886
N26210	FATLAX	-2624	12519	19243	9348	-2719	13428	19301	8592	-851
N72830	DTWFLX	5709	11642	17198	-153	5275	11111	19254	2867	2586
N12811	LAXOGG	6602	11110	17603	-109	6488	11116	19077	1472	1468
N14601	DCAHSV	5881	8490	17880	3508	5837	8946	19029	4244	692
N24715	BWISFO	7143	9590	15145	-1588	7049	9697	19028	2281	3776
N33608	PSPSEA	5706	9326	17688	2655	5633	9576	18988	3778	1050
N33785	CLEMKE	-2446	21033	18103	-484	-3053	18676	18983	3360	3237
N16217	LASMDW	-2600	22114	17634	-1879	-3438	20159	18958	2237	3278
N16806	DTWTPA	395	20308	15608	-5096	-368	16693	18923	2598	6931
N47332	JFKMDW	-295	15318	17468	2445	-444	15692	18892	3644	1049
N69803	ORDRSW	4028	13790	17431	-387	3898	13980	18844	965	1223
N12218	FLLORD	6047	13486	17160	-2374	5846	11789	18628	992	3165
N69351	MIATPA	-145	12670	17533	5008	-283	12739	18590	6134	987
N26232	MIATLH	-5881	24668	17775	-1010	-6751	21012	18257	3996	4137
N17644	BWIOAK	4222	14537	17245	-1513	3940	13880	18181	360	1592
N14662	OGGSFO	3774	15196	16608	-2361	3670	13922	18147	553	2811
N27722	CMHMDW	-6917	23302	17239	854	-7354	21229	18055	4180	2888
N19382	JAXLGA	8640	9072	15646	-2065	8571	8977	18051	502	2499

Table A.3: Airport Abbreviations

ABQ	ALBUQUERQUE INTL
ANC	ANCHORAGE INTL
AND	ANDERSON
ATL	ATLANTA HARTSFIELD
AUS	AUSTIN MUELLER
BDL	HARTFORD CT/SPRINGFIELD BRADLY
BNA	NASHVILLE METRO
BOS	BOSTON LOGAN
BUF	BUFFALO INTL
BWI	BALTIMORE INT'L
CHI	CHICAGO CHICAGO FSS
CHS	CHARLESTON MUNICIPAL
CIN	CARROLL
CLE	CLEVELAND HOPKINS
CMH	COLUMBUS INT'L
COS	COLORADO SPRINGS PETERSON
CVG	CINCINNATI CIN N.KNTY
DAL	DALLAS LOVE
DCA	WASHINGTON NATIONAL
DEN	DENVER STAPLETON
DET	DETROIT CITY
DFW	DALLAS INTL
DTW	DETROIT WAYNE CO
ELP	EL PASO INTL
EWK	NEW YORK NY/NEWARK NEWARK INTL
FAL	ROMA FALCON
FAT	FRESNO TERMINAL
FLL	FT. LAUDERDALE INTL
HAR	HARRISBURG SKYPORT
HNL	HONOLULU INTL
HOB	HOBBS LEA COUNTY
HOU	HOUSTON HOBBY
HSV	HUNTSVILLE/DECATUR HUNTSVILLE
IAH	HOUSTON INTERCONT
IND	INDIANAPOLIS INTL
INL	INTL FALLS
INT	GREENSBORO/H.PT/WIN-SALEM REYNOLDS
JAC	JACKSON HOLE
JAX	JACKSONVILLE INTL
JFK	NEW YORK NY/NEWARK KENNEDY
LAS	LAS VEGAS MCCARRAN
LAX	LOS ANGELES INTL
LGA	NEW YORK NY/NEWARK LA GUARDIA
LGB	LONG BEACH MUNICIPAL
LIH	LIHUE
LOG	LONGVIEW
MCO	ORLANDO INT'L
MDW	CHICAGO MIDWAY
MFE	MC ALLEN
MIA	MIAMI INT'L
MIT	SHAFTER KERN CTY
MKE	MILWAUKEE G MITCHELL
MRY	MONTEREY PENINSULA
MSP	MINNEAPOLIS INTL
MSY	NEW ORLEANS INTL
NEW	NEW ORLEANS LAKEFRONT
OAK	OAKLAND INTL
OMA	OMAHA EPPLEY
ORD	CHICAGO O'HARE
ORL	ORLANDO HERNDON
PBI	WEST PALM BEACH INTL
PDX	PORTLAND INTL
PHL	PHILADELPHIA PA/WILM'TON INTL
PHX	PHOENIX INTL
PSP	PALM SPRINGS
RDU	RALEIGH/DURHAM INTL
RIC	RICHMOND/WMBG INTL
RNO	RENO CANNON
RSW	FORT MYERS REGIONAL
SAN	SAN DIEGO LINDBERG
SAT	SAN ANTONIO INTL
SEA	SEATTLE/TACOMA SEA/TAC
SFO	SAN FRANCISCO INTL
SNA	SANTA ANA WAYNE INTL
STL	ST. LOUIS INTL
TLH	TALLAHASSEE MUNICIPAL
TPA	TAMPA TAMPA
TUS	TUCSON INTL
UMB	UMNAK ISLAND
VEL	VERNAL

Bibliography

- [1] United States Federal Aviation Administration. Airport capacity benchmark report, 2001. <http://www.faa.gov/events/benchmarks/download.htm>.
- [2] United States Federal Aviation Administration. Consolidated operations and delay analysis system. <http://www.apo.data.faa.gov/faacodasall.HTM>.
- [3] United States Federal Aviation Administration. Federal aviation regulations, section 121.471. http://www.faa.gov/avr/AFS/FARS/far_idx.htm.
- [4] Continental Airlines. Continental press release. http://www.continental.com/dash/build_dash.asp?press_2001-07-16-01.
- [5] ALPA. Agreement between northwest airlines, inc., and the air line pilots in the service of northwest airlines, inc, as represented by the air line pilots association, international. 1998-2002. Provided by the ALPA.
- [6] C. Barnhart and K. Talluri. Airline operations research. *Design and Operation of Civil and Environment Engineering Systems*, pages 435–469, 1997.
- [7] Victor Ilenda, Nathan Kleinman, and Stacy Hill. Spsa/simmod optimization of air traffic delay cost. *Proceedings of the American Control Conference*, June 1997.
- [8] Amadeo R. Odoni and Kerry M. Malone. The approximate network delay model. Received directly from Amadeo Odoni, August 1997.
- [9] Robert Simpson and Peter Belobaba. Air transportation economics, chapter 5. unpublished.

- [10] D.J. van Vuuren, P. Reithveld, and F.R. Bruinsma. Coping with unreliability in public transport chains : A case study for netherlands. *Transportation Research Part a*, 35:539–559, 2001.
- [11] William Vanderson. Improving aircraft departure time predictability. Master's thesis, Massachusetts Institute of technology, September 2000.