Probalilistic Analysis of Ground-Holding Strategies

by

Minakshi Sheel

B.S., University of Massachusetts, Amherst (1994)

Submitted to the
Sloan School of Management
in partial fulfillment of the requirements for the degree of

Master of Science in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 1997

Author ______________________________
Sloan School of Management
September 20, 1996

Certified by ______________________________
Amedeo R. Odoni
Professor, Department of Aeronautics and Astronautics and Department of Civil Engineering
Thesis Supervisor

Accepted by ______________________________
Thomas Magnanti
Co-Director, Operations Research Center

Probabilistic Analysis of Ground-Holding Strategies

by

Minakshi Sheel

Submitted to the Sloan School of Management
on September 20, 1996, in partial fulfillment of the
requirements for the degree of
Master of Science in Operations Research

Abstract

The Ground-Holding Policy Problem (GHPP) has become a matter of great interest in recent years because of the high cost incurred by aircraft suffering from delays. Ground-holding keeps a flight on the ground at the departure airport if it is known it will be unable to land at the arrival airport. The GHPP is determining how many flights should be held on the ground before take-off and for how long, in order to minimize the cost of delays. When the uncertainty associated with airport landing capacity is considered, the GHPP becomes complicated. A decision support system that incorporates this uncertainty, solves the GHPP quickly, and gives good results would be of great help to air traffic management.

The purpose of this thesis is to modify and analyze a probabilistic ground-holding algorithm by applying it to two common cases of capacity reduction. A graphical user interface was developed and sensitivity analysis was done on the algorithm, in order to see how it may be implemented in practice. The sensitivity analysis showed the algorithm was very sensitive to the number of probabilistic capacity scenarios used and to the cost ratio of air delay to ground delay. The algorithm was not particularly sensitive to the number of periods that the time horizon was divided into. In terms of cost savings, a ground-holding policy was the most beneficial when demand greatly exceeded airport capacity. When compared to other air traffic flow strategies, the ground-holding algorithm performed the best and was the most consistent under various situations. The algorithm can solve large problems quickly and efficiently on a personal computer.

Thesis Supervisor: Amedeo R. Odoni
Title: Professor, Department of Aeronautics and Astronautics and Department of Civil Engineering

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Air transportation networks today experience high costs from aircraft delays. Air traffic flow management attempts to schedule air traffic demand against airport and airspace capacity, on a daily basis, in a manner that minimizes the cost of these delays. Ground-holding is used as a strategy to manage the flow of air traffic. The primary reason for ground-holding is to enhance the safety and economy. It is safer and less expensive to hold a flight on the ground at its airport of departure, rather than have the flight delayed in the air, if it is known that there will be delays at its destination airport. The Ground-Holding Policy Problem (GHPP) is concerned with deciding which flights should be delayed on the ground before take-off and for how long, in order to minimize the cost of delays.

The GHPP becomes complicated when the probabilistic nature of airport capacities is introduced. Airport capacities for arrivals, as well as departures, are quite variable. Weather conditions, airport activity, mechanical problems or human factors can have a significant effect on capacity, though weather usually is the factor most responsible for creating uncertainty in capacity predictions.

In this thesis, we look at the single airport static GHPP. We are interested in a model that incorporates the uncertainty in airport landing capacity. This is needed in order to have an accurate representation of reality and to give good solutions, in terms of delay times and costs. It is also important that the model be able to be easily and usefully implemented in practice. Airports need a decision support system that can solve large problems in real-time and give the solutions in a simple, easy to read manner. One way of accomplishing this would be a computer program that solves the problem quickly, is user friendly, and gives answers in graphics form. In the following chapters, we analyze a decision support system by applying a GHPP model to two cases of capacity reduction that occur at many airports in the U.S.

The outline of this thesis is as follows. Chapter 2 reviews some relevant work that has already been done on air traffic flow management, mainly Richetta's algorithm for the single airport static ground-holding problem with probabilistic capacity scenarios. In Chapter 3, we show how Richetta's model can be applied to two common cases of capacity reduction, usually caused by the arrival of some kind of weather condition. We show how this simplifies the formulation of the model, in terms of the number of variables and constraints. The computer implementation of the model and its connection to a graphical user interface (GUI) are then described. Chapter 4 discusses five areas of experimentation that are performed in order to better understand Richetta's algorithm and how it can be applied in practice. The first set of experiments are to see how an instance of capacity reduction affects an airport differently at different times of the day. We do these tests on three airports with different demand profiles in order to observe how various airports are affected differently by a period of capacity reduction. The second and third areas of experimentation test the algorithm's sensitivity to the number of capacity scenarios used and to the number of time intervals the period of interest is divided into. In these two sections, we are interested in the impact of the number of capacity scenarios and time intervals on the results and solution time for the problem. The fourth set of experiments show how the ratio of the cost of air delay per time period to the cost of ground delay per time period can determine what kind of ground-holding policy will be implemented. In the fifth area, we see how our algorithm performs in comparison to other air traffic flow strategies. Finally in Chapter 5, conclusions are drawn. The ground-holding algorithm can solve large problems quickly and performs well against other strategies, especially in situations where demand greatly exceeds capacity. Because of the algorithm's sensitivity to the number of capacity scenarios used and to the cost ratio, these values should be determined with care. In this last chapter, we also delve into potential topics for future research.

Chapter 2

Air Traffic Flow Management

2.1 Literature Review

Air traffic management operations can be subdivided into air traffic control, which tries to avoid "conflicts" and collisions between aircraft, and air traffic flow management, which tries to optimize the flow of aircraft. Research done on the latter subject is the focus of this chapter.

The two main problems in air traffic flow management are the ground-holding problem and the problem of deciding how to redistribute flows in the air sectors. These problems arise when demand exceeds the capacity at airports or air sectors.

Because of the stochastic nature of airport capacity, probabilistic models represent the GHPP better than deterministic ones. A deterministic model has only one capacity profile, whereas a probabilistic one can have multiple capacity profiles, each with an associated probability. The GHPP can also be divided into the Single-Airport GHPP and the Multi-Airport GHPP.

In the single-airport case, flights are scheduled to land at one "arrivals" airport, which is the only capacitated element of the system under consideration. The model can be either static or dynamic. Static means that all decisions for ground-holding are made once, at the beginning of the time horizon, while in the dynamic model, ground-holding decisions are continually revised during the day, as new information about capacity is received. A deterministic, static model is described in Terrab (1990) and in Terrab and Odoni (1993). A probabilistic, static model is explained in Richetta (1991) and Richetta and Odoni (1993). (This thesis explores this model and a detailed description is given in the following subsection.) Richetta and Odoni (1994) go on to develop a probabilistic, dynamic model which is a more general model of the single-airport GHPP.

In the multi-airport GHPP case, a network of airports is examined. This is done in order to model situations in which there are connecting flights that may propagate delays throughout the network of airports. This model has capacity constraints for arrivals and departures, as well as connectivity constraints. There are four optimization models and two heuristics, which essentially cover this topic. All existing models are deterministic. The four optimization models are discussed, respectively, in Vranas (1992) and Vranas, Bertsimas, and Odoni (1994a, 1994b); Andreatta and Brunetta (1995); Bertsimas and Stock (1994); and Terrab and Paulose (1993). It is believed that the Bertsimas and Stock model performs best in most cases (Andreatta and Brunetta (1995)). The two heuristics are based on "priority rules", which are used to assign ground-holding times to sequences of aircraft at each airport. These two heuristics are by Andreatta, Brunetta, and Guastalla (1994) and by Navazio and Romanin-Jacur (1995).

Another area of air traffic flow management research is concerned with redistributing flows in airspace by rerouting flights. This problem has been analyzed by using an origin-destination network that models airports and enroute air traffic sectors. Existing models in this area are deterministic. An important optimization model is the one by Bertsimas and Stock (1994). Computationally, this has been the most successful model so far. Two other optimization models are the Time Assignment Model [Boyd, Burlingame, and Lindsay (1993, 1994, 1995)] and the Space-Time Network [Helme (1992, 1994)]. There are also two heuristics, the Multiple Airport Scheduler [Epstein, Futer, and Medvedovsky (1992)] and the Computer Assisted Slot Allocation Algorithm [Philipp and Gainche (1994)].

This chapter focuses on how to optimally solve the single-airport static GHPP, as described by Richetta. The model is a probabilistic one that incorporates various airport landing capacity profiles, each with a particular probability, into the GHPP model.

2.1.1 Richetta

The model, and the ensuing algorithm, described in "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control" by Richetta and Odoni is the basis for the work done in this thesis. Therefore, it is important to give a detailed description of this model.

The model was developed in order to incorporate into GHP decision-making the uncertainty in airport capacity. Because of this uncertainty, there will be times when available landing capacity is wasted, even with an optimal ground-holding policy. This is because air traffic managers tend to be "conservative" if expecting bad weather. If they are too "liberal", then there will be instances where aircraft suffer excessive delays in the air because there is not enough landing capacity. When considering the trade-off between a conservative policy, that may assign too much ground-holding and waste available capacity, and a liberal policy, that may lead to more expensive air delays, air traffic management must weigh the cost of holding an aircraft on the ground versus the cost of having an aircraft absorb the delay in the air. Thus, a good ground-holding policy must consider the probabilities of the various possible values of airport capacity, as well as the cost ratio of ground and air delay.

Richetta developed a static model for the multi-period single-airport GHPP. It can be solved on a personal computer using stochastic linear programming with one stage. The model allows for different probabilistic capacity scenarios and different aircraft classes. The GHPP is solved on an aggregate level. This means the solution gives information on how many aircraft should be delayed and for how long, but does not say which individual aircraft should be delayed.

Fig. 2-1 : Simplified Model

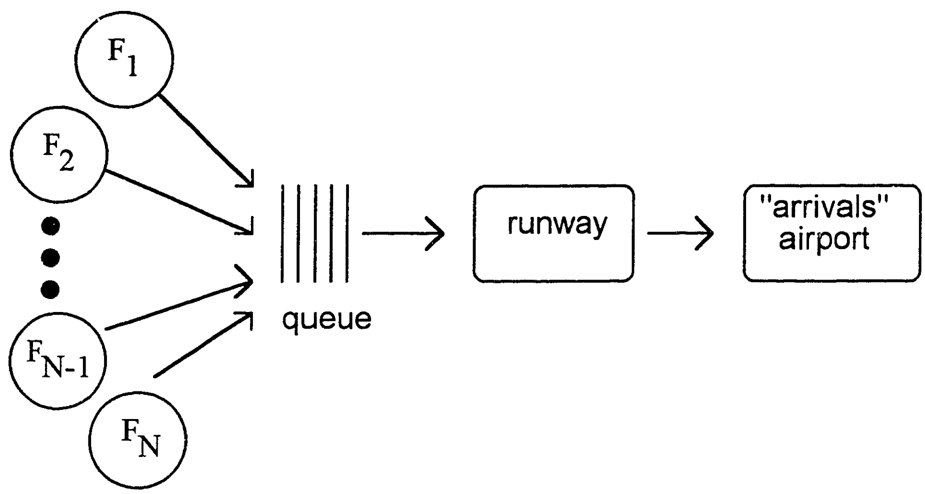

A representation of the simplified model can be seen in Fig. 2-1. The key ideas of this model are:

- There is only one "arrivals" airport. N aircraft ($F_1,\ldots,F_N$) leave the various departure airports and form a queue for the runway at the "arrivals" airport.
- The "arrivals" airport is considered to be the only element in the model with restricted capacity.
- The departure time and flight length of all aircraft are deterministic and known at the beginning of the time interval.
- The time interval, [0,B], is divided into T equal parts. The earliest departure is at time 0 and the latest arrival is at B.
- Because this is a static model, the assumption is that Q different capacity scenarios and their associated probabilities, P_q, are given at time 0. The capacity at period T+1 equals N, the total demand of aircraft for the "arrivals" airport. This ensures that all aircraft will be able to land within T+1 periods. Since demand decreases toward the end of the day, it is plausible to assume all flights will land by the end of the day.
- The cost functions for ground and air delay are known for each aircraft.

Notation is defined as follows:

M_{qi} : The capacity for period i under scenario q. The capacity for period T+1=N. (q=1,...,Q; i=1,...,T+1)

N_{ki} : The number of aircraft of class k scheduled to land at the "arrivals" airport in period i. (k=1,...,K; i=1,...,T)

X_{qkij} : The decision variables showing how many aircraft of class k scheduled to land at the "arrivals" airport in period i, but are rescheduled to arrive in period j under capacity scenario q, with a ground delay of (j-i) time periods. (q=1,...,Q; k=1,...,K; i=1,...,T; $i \leq j \leq T+1$)

$C_g(k,i)$: The cost of delaying an aircraft of class k on the ground for i periods. (k=1,...,K; i=1,...,T-1)

c_a : The cost of delaying an aircraft in the air for one period. This is assumed to be a constant because of the FCFS decision for aircraft waiting to land.

W_{qi} : The number of aircraft that cannot land at the "arrivals" airport in period i under capacity scenario q. (q=1,...,Q; i=1,...,T)

The final formulation for the static GHPP is reached after four stages. The first stage gives an integer linear programming problem for one capacity scenario and one class of aircraft. The second stage transforms this formulation into a min. cost flow problem. In the third stage, the probabilities are introduced and the min. cost flow problem becomes a distribution problem. Finally, in the fourth stage, the formulation changes when a set of coupling constraints are added.

In the first stage, in order to simplify the model, it is assumed that there is only one capacity scenario with probability one and only one type of aircraft class. The formulation is an integer linear programming problem and is as follows:

$$\text{Min} \sum_{i=1}^{T} \sum_{j=i+1}^{T+1} C_g\,(j-i)\, X_{qij} + \sum_{i=1}^{T} W_{qi}\, c_a$$

S.T.

$$\text{(i)} \quad \sum_{j=i}^{T+1} X_{qij} = N_i \qquad i=1,\dots,T$$

$$\text{(ii)} \quad W_{qi} \geq \sum_{j=1}^{i} X_{qji} + W_{qi-1} - M_{qi} \qquad i=1,\dots,T+1;\ W_{q0} = W_{qT+1} = 0$$

$$\text{(iii)} \quad X_{qij}\,,\ W_{qi} \geq 0 \text{ and integer}$$

The objective function is a linear cost function that minimizes the cost of ground and air delays. Constraints (i) ensure that all aircraft scheduled to land in period i will indeed land before the end of the day. This means that an aircraft meant to arrive during period i will land in a time somewhere between periods i and T+1. The second set of constraints comes from the flow balance at the end of each period for the "arrivals" airport.

The second stage converts the above formulation into a minimum cost flow problem. A network picture of this can be seen in Fig. 2-2. The left column of nodes are the supply nodes with supply N_i and the center column of nodes are the demand nodes with demand M_i. The cost on the arcs going from the supply nodes to the demand nodes is the cost of ground-hold delay and the flow represents the rescheduling of aircraft. The arcs that go from one demand node to the subsequent demand node have a cost equal to the cost of airborne delay and a flow equal to the number of aircraft that incur air delay . The node in the right column is a surplus node and all arcs leaving this node to go to the demand nodes have zero cost and a flow of surplus capacity.

Fig. 2-2: Minimum Cost Flow Problem

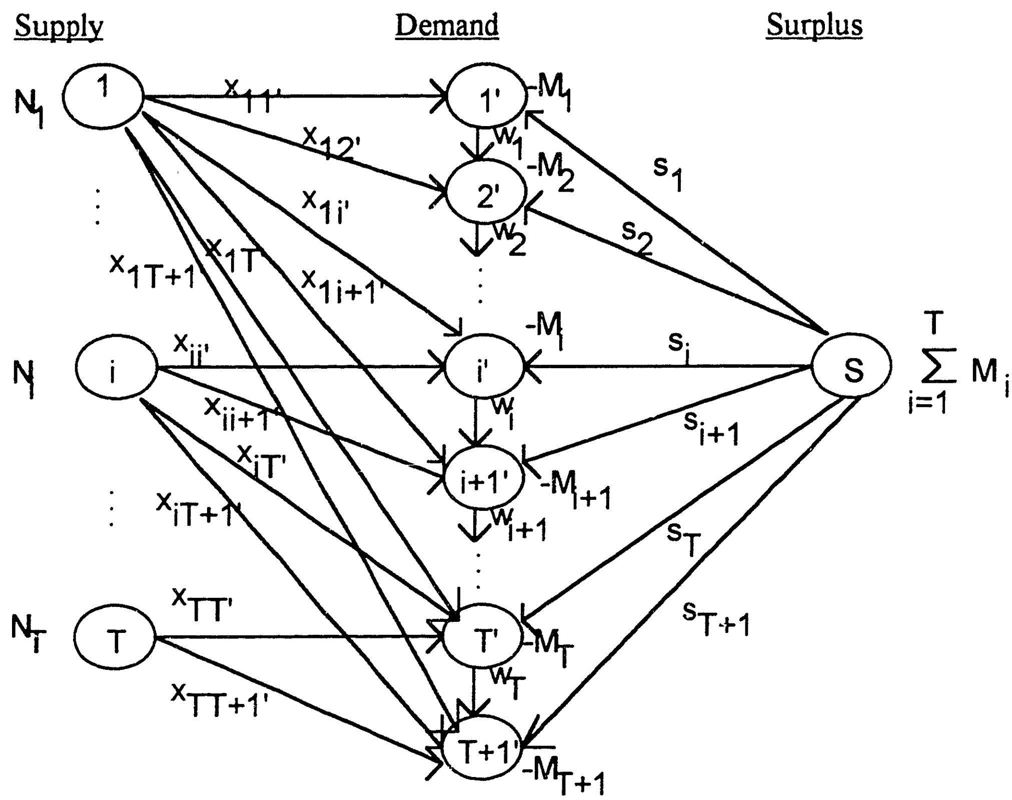

The formulation for this network then becomes:

$$\text{Min} \sum_{i=1}^{T} \sum_{j=i+1}^{T+1} C_g\,(j\text{-}i)\, X_{qij} + \sum_{i=1}^{T} W_{qi}\, c_a$$

S.T.

(i) $\sum_{j=i}^{T+1} X_{qij} = N_i \qquad i=1,\ldots,T$

(ii) $\sum_{i=1}^{T+1} S_{qi} = \sum_{i=1}^{T} M_{qi}$

(iii) $W_{qi} - (W_{qi-1} + \sum_{j=1}^{i} X_{qji} + S_{qi}) = - M_{qi} \qquad i=1,\ldots,T+1;\ W_{q0} = W_{qT+1} = 0$

(iv) X_{qij}, W_{qi}, $S_{qi} \geq 0$

The objective function has not changed from that of the first stage. Constraints (i) and (ii) correspond to the supply nodes while constraints (iii) correspond to the demand nodes.

The integrality constraints have been relaxed, though. This is possible because of the total unimodularity of the constraint matrix and the fact that the supply and demands are integer.

Uncertainty is introduced in the third stage and the minimum cost flow problem becomes a distribution problem. To reiterate, there is only one class of aircraft and the different probabilistic capacity scenarios for all time periods are known at time zero. The formulation now becomes:

$$\text{Min} \quad \sum_{q=1}^{Q} P_q \left(\sum_{i=1}^{T} \sum_{j=i+1}^{T+1} C_g\,(j-i)\, X_{qij} + c_a \sum_{i=1}^{T} W_{qi} \right)$$

S.T. for each scenario q=1,...,Q:

(i) $\sum_{j=i}^{T+1} X_{qij} = N_i \qquad i=1,\ldots,T$

(ii) $\sum_{i=1}^{T+1} S_{qi} = \sum_{i=1}^{T} M_{qi}$

(iii) $W_{qi} - (W_{qi-1} + \sum_{j=1}^{i} X_{qji} + S_{qi}) = - M_{qi} \qquad i=1,\ldots,T+1;\ W_{q0} = W_{qT+1} = 0$

(iv) $X_{qij}, W_{qi}, S_{qi} \geq 0$

Notice that the above constraints are identical to those in stage two. This means that the distribution problem can be solved by solving Q separate min. cost flow problems. The optimal solution will give a ground-holding policy for each scenario, but in practice, only one policy can be followed because a policy must be made before the actual capacity is known. Therefore, in stage four, a set of constraints that sets X_{qij} equal for all q must be added and is as follows:

$$X_{1ij} = X_{2ij} = X_{3ij} = \ldots\ldots\ldots = X_{Tij} \qquad i=1,\ldots,T;\ i \leq j \leq T+1$$

This coupling constraint now causes total unimodularity to be lost, and it is still unknown whether unimodularity indeed exists in this problem. While total unimodularity is a sufficient condition for obtaining integer solutions to a relaxation of an integer linear program, unimodularity is necessary. However, Richetta proceeds to relax the integrality constraints because he could not come up with a counterexample of non-integer solutions and even if solutions were to be non-integer, they could be rounded without ruining the goodness of a solution. It should be noted here that we did about a few hundred runs of the formulation and did not once get a non-integer solution.

The coupling constraints can be removed via substitution into the distribution problem formulation. The new constraint matrix makes it impossible to solve the problem with fast decomposition techniques, but there will now be fewer constraints and variables, making it possible to solve quickly on a personal computer. Now, all that is needed is to rewrite the formulation so that it includes K classes of aircraft. The final formulation thus looks like:

$$\text{Min} \quad \sum_{k=1}^{K} \sum_{i=1}^{T} \sum_{j=i+1}^{T+1} C_g(k,j-i)\, X_{kij} + c_a\Big(\sum_{q=1}^{Q} P_q \sum_{i=1}^{T} W_{qi}\Big)$$

S.T.

$$\text{(i)} \quad \sum_{j=i}^{T+1} X_{kij} = N_{ki} \qquad k=1,\dots,K;\ i=1,\dots,T$$

$$\text{(ii)} \quad \sum_{i=1}^{T+1} S_{qi} = \sum_{i=1}^{T} M_{qi} \qquad q=1,\dots,Q$$

$$\text{(iii)} \quad W_{qi} - W_{qi-1} - \sum_{k=1}^{K} \sum_{j=1}^{i} X_{kji} - S_{qi} = -M_{qi} \qquad q=1,\dots,Q;\ i=1,\dots,T+1;\ W_{q0}=W_{qT+1}=0$$

(iv) X_{kij}, W_{qi}, $S_{qi} \geq 0$ and integer

To see how "good" his formulation was, Richetta compared its performance to that of two other algorithms, the deterministic and the passive. The deterministic algorithm chooses the capacity scenario with the highest probability as the only scenario and proceeds to solve the static GHPP by assigning available capacity as FCFS, with all delays being ground-hold ones. The passive algorithm has a planning strategy that allows all aircraft to take off according to schedule and any delays will be airborne delays. There are multiple probabilistic capacity scenarios and these are used to figure out the air delays. The passive algorithm will give the minimum amount of delay possible.

Richetta ran his algorithm for a single class of aircraft and evaluated its performance against these two algorithms for different demand and capacity cases. The static algorithm had minimum cost, with the deterministic as second cheapest, and the passive with the highest cost. However, as the cost of air delays increases, the advantage of static diminishes in respect to deterministic. This is because deterministic has the lowest amount of air delay. But generally, static will give minimum cost. Another important result was that the total expected delays, when using the static algorithm, was within 15% of the minimum expected delays, the amount given by the passive algorithm. Though static has more delay time than the minimum required, most of its delay are ground-holding delays which are cheaper than airborne delays. Thus, Richetta's algorithm appears to be a very efficient algorithm that can be used to solve problems of sizes comparable to actual situations quickly on a personal computer.

Chapter 3

Development of the Static Ground-Holding Problem

In this chapter, our simplified version of Richetta's model for the static ground-holding problem is introduced and explored. The simplified model deals with a single class of aircraft, instead of several classes. The formulation of this model and the changes in it, due to the simplification, are described in the first section of this chapter. The second section shows two common cases of airport capacity reduction that this model can be applied to. In the third section, the computer implementation of this formulation is described.

3.1 Formulation of the Model

As reviewed in Chapter 2, Richetta's Algorithm for the static ground-holding problem can be solved for K classes of aircraft. However, common practice in the United States today assigns available airport capacity for arrivals of aircraft on a first-come first-serve basis. This implies that the need to distinguish between type of aircraft, in order to determine priority for a flight in the "arrivals" queue, is superfluous. Thus, having only one class of aircraft, when determining a ground-holding policy, would comply with current practice in the U.S.

Treating all arriving aircraft equally will simplify the ground-holding cost function, $C_g(k,i)$, which is the cost of delaying an aircraft of class k on the ground for i periods. This will now depend solely on the number of time periods an aircraft is held on the ground. We chose, though, to use a linear cost function and treat the ground-holding cost as a constant, c_g, per time period. This decision allows the effect of the ratio of air-hold cost to ground-hold cost to be analyzed. Different ratios can affect a policy to be either conservative or liberal. A high ratio of air to ground cost will result in a great deal of ground-holding because it is far less expensive than air delays (a "conservative" strategy). Making the air cost slightly less than the ground cost will ensure a passive ground-holding strategy. This means all aircraft will leave the airports of origin according to schedule and all delays will be absorbed in the air. Thus, varying c_a and c_g allows us to see how the problem is solved for different ground-holding strategies.

Incorporating these changes into the final formulation from Chapter 2 (page 18), we obtain a simplified version, that is as follows:

$$\text{Min} \quad c_g\left(\sum_{i=1}^{T} \sum_{j=i+1}^{T+1} (j-i)\, X_{ij}\right) + c_a\left(\sum_{q=1}^{Q} P_q \sum_{i=1}^{T} W_{qi}\right)$$

S.T.

$$\text{(i)} \quad \sum_{j=i}^{T+1} X_{ij} = N_i \qquad i=1,\ldots,T$$

$$\text{(ii)} \quad \sum_{i=1}^{T+1} S_{qi} = \sum_{i=1}^{T} M_{qi} \qquad q=1,\ldots,Q$$

$$\text{(iii)} \quad W_{qi} - W_{qi-1} - \sum_{j=1}^{i} X_{ji} - S_{qi} = -M_{qi} \qquad q=1,\ldots,Q;\ i=1,\ldots,T+1;\ W_{q0}=W_{qT+1}=0$$

(iv) X_{ij}, W_{qi}, $S_{qi} \geq 0$ and integer

The number of variables and constraints in the above formulation has been reduced significantly. In the formulation where K aircraft classes are considered, there are $(KT+(T+2)Q)$ constraints and $(K((T^2+3T)/2)+(2T+1)Q)$ variables, where T is the number of time periods and Q is the number of probabilistic capacity scenarios. These values were obtained in the following manner:

Number of Constraints

- For the set of constraints (i), there are KT constraints.
- For the set of constraints (ii), there are Q constraints.
- For the set of constraints (iii), there are Q(T+1) constraints.
- Summing the above gives KT+(T+2)Q constraints.

Number of Variables

- For the variables X_{kij}, where k=1,...,K, i=1,...,T, and $i \leq j \leq T+1$, there are

$$K\sum_{i=2}^{T+1} i = K((T^2+3T)/2) \text{ variables.}$$

- For the variables S_{qi}, where q=1,...,Q and i=1,...,T+1, there are Q(T+1) variables.
- For the variables W_{qi}, where q=1,...,Q and i=1,...,T, there are QT variables.
- Summing the above gives $K((T^2+3T)/2)+(2T+1)Q$ variables.

In the simplified version, K=1 and it can easily be seen that the number of constraints and variables are now, respectively, (T+(T+2)Q) and $((T^2+3T)/2)+(2T+1)Q$. Thus, the number of constraints and variables is more dependent on T.

3.2 Application of the Model

In this thesis, we look at a single airport's schedule of arrivals and the landing capacity for these arrivals over a time horizon of one day. The schedule of arrivals is deterministic and the capacity is probabilistic. The ground-holding model described in this thesis was developed so that the probabilistic nature of airport capacities could be incorporated into decision making on ground-holding. The uncertainty associated with the landing capacity is often a result of adverse weather conditions. This section describes the two common cases of capacity reduction that will be analyzed in Chapter 4. These two cases are when there is uncertainty in the occurrence time of capacity reduction and when there is uncertainty in the amount of capacity reduction.

3.2.1 Uncertainty in Time

This is the case where the amount of capacity reduction and the duration of the reduction are approximately known, but the time of occurrence is not. A condition that could result in this type of situation is a blizzard passing through the area where a major airport is located. When this occurs, if the airport does not shut down completely, the available runway capacity is reduced significantly. It may be difficult to pinpoint the exact arrival time of the blizzard. Current weather forecasting technology can follow the path of the blizzard and be able to give a good approximation of its duration, but the arrival time is usually given in a range. For example, a typical weather forecast may say that the blizzard is expected to arrive between 9:00 AM and noon and will last about eight hours. Fig. 3-1 shows a graph of this case of capacity reduction, where the uncertainty is in the start time of the reduction.

Fig.3-1: Uncertainty in Time

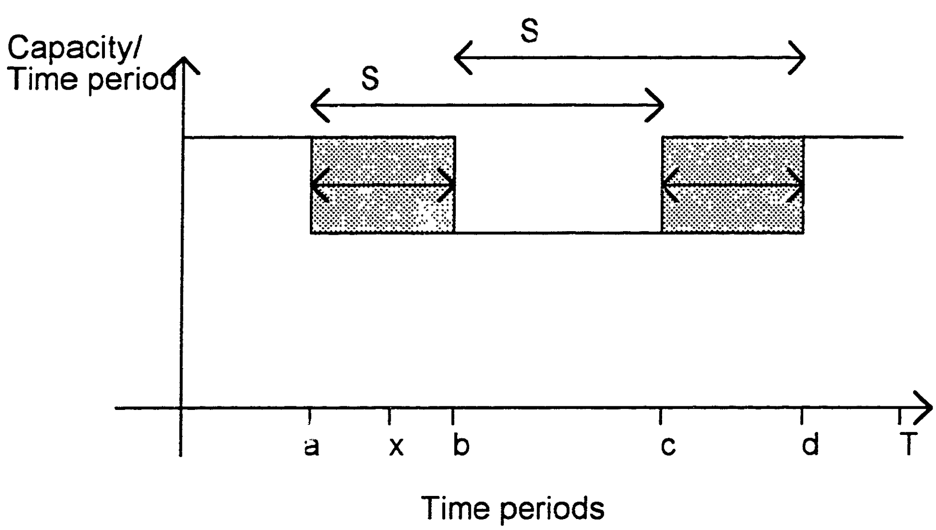

In Fig. 3-1, some kind of weather condition is expected to arrive between a and b, thereby reducing capacity for a length of time, S. It should be noted that only the arrival time is uncertain. No matter when the capacity reduction occurs, it will last for S time units. If a storm arrives at time a, capacity reduction will last until time c and if the storm arrives at time b, it will last until time d. Similarly, if the storm arrives for some time x, between a and b, capacity reduction will last until time x+S. There is a probability distribution for this storm arrival time. The distribution could be uniform, triangular, normal, one-tailed, etc. Because the time axis is in discrete time intervals, a probability could be associated with each time increment between a and b, including a and b. If we assume time increments are one time unit apart, then there would be (b-a+1) probabilities, each one corresponding to a capacity scenario. (If a=3 and b=5, then a storm could arrive at 3, 4, or 5, giving three possibilities, 5-3+1=3.) Capacity Scenario 1 would then have the capacity drop from a until c, when it would go back up. Capacity Scenario 2 would drop from (a+1) until (c+1) and then go back up. This reasoning continues through Capacity Scenario (b-a+1), where the capacity drops from b until d, where it then goes back up. If the distribution was uniform, then each scenario would have a probability associated with it equal to 1/(b-a+1). Another possibility is that instead of a probability profile for every time increment between a and b, we could have three profiles. Under Capacity Scenario 1 the drop in capacity would begin at a (probability 1/3), under Scenario 2 it would begin at x (probability 1/3), and under Scenario 3 it would begin at b (probability 1/3). A question that arises is whether, for the uniform distribution (or any distribution), we need (b-a+1) capacity scenarios or whether fewer scenarios might be a sufficient approximation.

In order to further clarify the above situation, we will use an example. Let us say that at the beginning of the day, we are given a weather forecast that predicts a storm is going to arrive at Logan Airport some time between 10 AM and noon and will last roughly four hours. During the storm, airport capacity will be reduced from 60 arrivals per hour to 36 arrivals per hour. Also, the storm is equally likely to begin at any time between 10 AM and noon.

In this example we are looking at the capacity for a 17 hr. day that begins at 7 AM and ends at midnight. Because demand between midnight and 7 AM is small, we will assume that all flights are able to land in this time frame. Each hour is divided into 15 minute periods, giving the total number of periods, T, equal to 68 (17 hrs.*4=68). For this example, the highest level of capacity is 60 arrivals/hr. and the storm will reduce the capacity to 36 arrivals/hr. Because each hour has been divided into 15 minute intervals, these two capacities, respectively become, 15 arrivals/15 min. and 9 arrivals/ 15 min. Fig. 3-2 illustrates this example. The values on the x-axis, time, and the values on the y-axis, capacity, are discrete integers.

Fig. 3-2: Example of Uncertainty in Time (not drawn to scale)

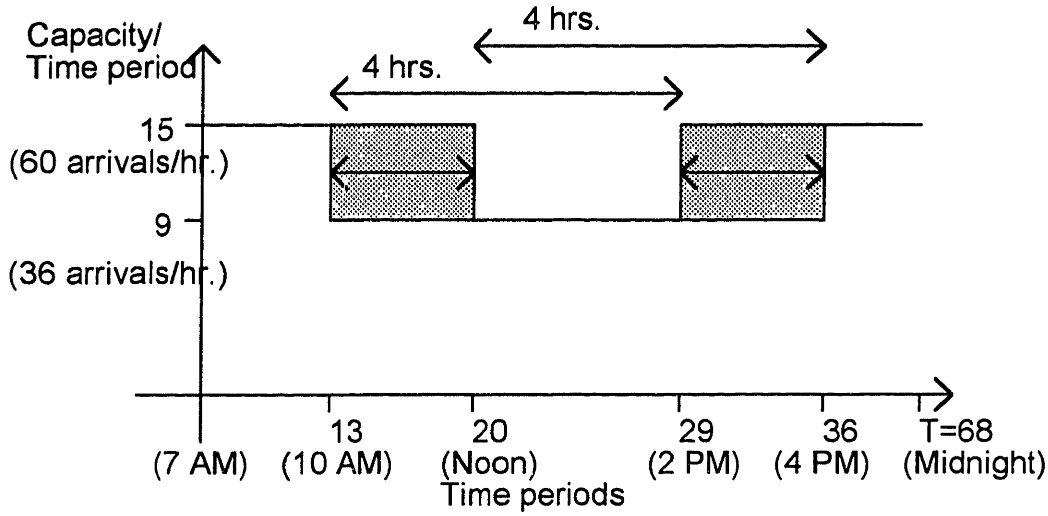

Therefore, in this example, capacity will be at 15 arrivals/15 min. up to at least period 13 (10 AM), then at some time between time periods 13 and 20, capacity will drop down to 9 per 15 min., for 16 time periods (4 hrs.) and then capacity will go back up to 15 arrivals/15 min. Thus, the range of time intervals that capacity will be reduced is [(13-29), (14-30),..., (19-35), (20-36)] and the reduction is equally likely to occur at any one of these time intervals. There would then be eight capacity scenarios, as shown in the shadowed range, and each one would have a probability of .125. One of the issues we shall investigate is whether these 8 scenarios can be replaced, without much loss in accuracy, by fewer scenarios, e.g., by 4 which begin at 30 minute intervals instead of 15 minute ones.

3.2.2 Uncertainty in Capacity

The second type of capacity reduction that will be examined is the case where the time of occurrence and the duration of the capacity reduction are approximately known, but the amount of reduction is not. An instance of this case may be an airport that often has a great deal of fog, like San Francisco. When fog occurs in the early part of the day, it may have begun in predawn and it will be dispersed soon after the sun rises, if it is going to be a sunny day. Therefore, when the fog begins, the length of time it will last is roughly known, but how badly this will affect the airport capacity is uncertain. Fig. 3-3 shows this type of capacity reduction.

Fig. 3-3: Uncertainty in Capacity

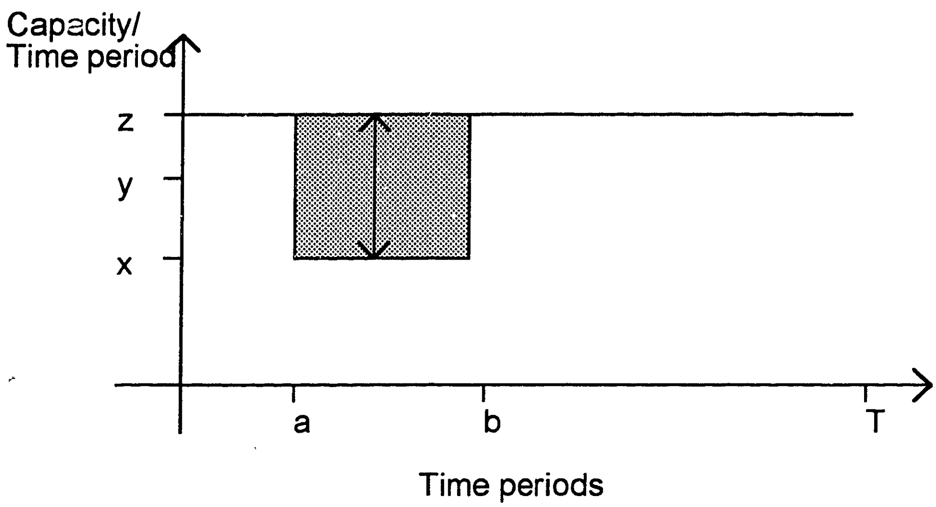

In Fig. 3-3, capacity reduction will occur at time a and last until time b. These are assumed to be known values. The capacity per time period is the factor that has uncertainty associated with it. For the time period a to b, capacity will be somewhere between the values of x and z, but the exact value is unknown. From time zero to a, and then from time b to T, airport capacity will be z. It is only between a and b, that the capacity may drop. For the amount of capacity reduction, there is a probability distribution. Every integer value between x and z has an associated probability. If the distribution is uniform, then the values x, x+1,..., y,..., z-1, z each would have a probability of 1/(z-x+1).

Just as in the previous subsection, we will use an example to make things more clear. Fog is expected to roll in after midnight and is not expected to disappear until around 10 AM. Depending on how dense the fog is, airport landing capacity will be somewhere between 36 arrivals per hr. and 60 arrivals per hr. The capacity is equally likely to be anywhere between these two values. Breaking one hour into four 15 minute periods and looking at a time horizon between 7 AM and midnight, T will equal 68. Airport capacities of 36 arrivals/hr. and 60 arrivals/hr. become respectively, 9 arrivals/15 min. and 15 arrivals/15 min.. This is shown in Fig. 3-4.

Fig. 3-4: Example of Uncertainty in Capacity

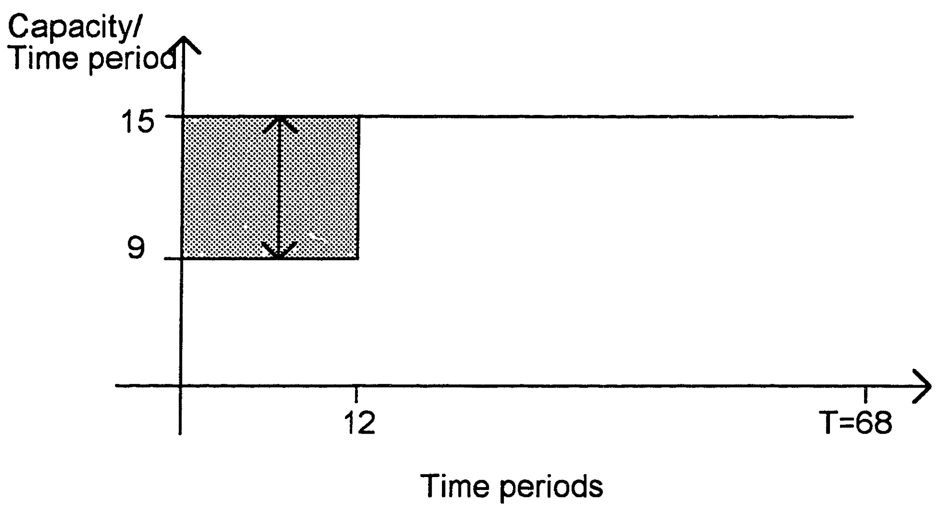

In this example, for the time periods 1 to 12, the capacity will be somewhere between 9 and 15 arrivals/15 min. For the time periods 13 to 68, capacity will be at the maximum of 15 arrivals/15 min. There will be seven capacity scenarios, each with a probability of 1/7. Capacity Scenario 1 would have a capacity of 15 for all time periods. Capacity Scenario 2 would have a capacity of 14 for time periods 1-12 and then the capacity would be 15 for time periods 13-68. This continues until capacity scenario seven, when the capacity is 9 for time periods 1-12 and then is 15 from time period 13 on.

3.3 Computer Implementation

All experiments were done on Sun Sparc 10 stations. To solve the static ground-holding formulation, the LP optimizer CPLEX was used. Either the C language or the General Algebraic Modeling System (GAMS) package is needed to put the formulation in a form that CPLEX can read and then solve. First, GAMS was tried because it is rather straight-forward and relatively easy to learn. With GAMS, though, we ran into memory problems. We, therefore, used the C language in order to take full advantage of the callable libraries of CPLEX.

The C program was linked to a graphical user interface (GUI), which was coded in X-Windows using Motif to display inputs and outputs. Inputs can thus be changed easily and the program rerun. The inputs and outputs are as follows:

Inputs

- c_a : cost of delaying an aircraft in the air for one period
- c_g : cost of delaying an aircraft on the ground for one period
- Q : number of capacity scenarios
- P_q : probability associated with capacity scenario q
- N_i : the original schedule of flights for periods i=1,...,T
- M_{qi} : for scenario q, the capacity for periods i=1,...,T+1, where $M_{qT+1}=\sum_{i=1}^{T} N_i$

Outputs

- The new schedule of flights is: $NFS_j=\sum_{i=1}^{j} X_{ij}$ where i=1,...,T; $i \le j \le T+1$
- The number of flights unable to land during period i under capacity scenario q is: W_{qi} where q=1,..,Q; i=1,...,T
- The total expected air delay is: Air delay$=\sum_{q=1}^{Q}(p_q \sum_{i=1}^{T} W_{qi})$
- The total ground-holding delay is: Ground delay$=\sum_{i=1}^{T}\sum_{j=i+1}^{T+1}(j-i)X_{ij}$
- The total delay: Total delay= Total Ground delay + Total Air delay
- The cost of total delay in the air: Air delay cost= c_a(Air delay)
- The cost of total ground-holding delay: Ground delay cost= c_g(Ground delay)
- The total cost is: Total delay cost= Ground delay cost + Air delay cost

When the program for the GUI is called, a window appears on the screen as shown in Fig. 3-5 (see end of chapter). The top two buttons on the right hand side are, respectively, c_a and c_g. The third button is for Q (maximum value is 10) and it is followed by p_q (q=1,...,10). These all can be edited by clicking on the button of choice and typing in the desired value. To load an input file, we go to the Load button and click on an input file of choice. Plots of the original schedule of flights (demand) and the capacity scenarios will appear on the graph of the window. To run the program, we click on the Run button. When the solution is found, two more windows appear on the screen. See Fig. 3-6 and Fig. 3-7 at the end of chapter. The new schedule of flights (optimal solution) appears on the same graph, in the main window, as the demand and capacity. Fig. 3-6, one of the windows that pops up, gives a table of the Air delay, Ground delay, Total delay, Air delay cost, Ground delay cost, and Total delay cost. The other window, Fig. 3-7 shows W_{qi}, the number of flights that will be delayed in the air, under this optimal solution, if capacity scenario q were to actually occur.

Now, if we choose to edit the input, we do not have to go to the input file. It can all be done in the main window. The costs, Q, and p_i can be edited as described earlier, and then the file is reloaded. To edit the demand or one of the capacity scenarios, we go to the Edit button and click on what we wish to edit. Whatever is chosen appears alone on the main graph. The mouse is used to make desired changes. For example, suppose we wish to edit a capacity scenario which currently has a capacity of 12 for all time periods and suppose we want the capacity for time periods 20-40 to be 10. We would then click on the point (20,10) on the graph with the left button of the mouse, click on the point (40,10) with the right button, and then click the middle button. We could continue in this manner to do as many changes as desired. Under the Quit button, an (X,Y) coordinate is shown. This tells us where on the graph is the cursor. After all changes are done, we go back to the Edit button and click on Commit. This incorporates the change. We could proceed in this way to edit the demand or other capacity scenarios. Afterwards, we click on the Run button to rerun the program. The Clear button eliminates the optimal solution from the main graph. The Clear All button clears everything from the main graph. If we wish to print on a printer, we go to the Print button and then click on the window we wish to print out. The Print to file button is used the same way, except we designate a file to print the window to.

The GUI was developed in order to easily manipulate the input data and to view the solution to the linear program in an easy to read manner.

Fig. 3-5 Main Gui Window

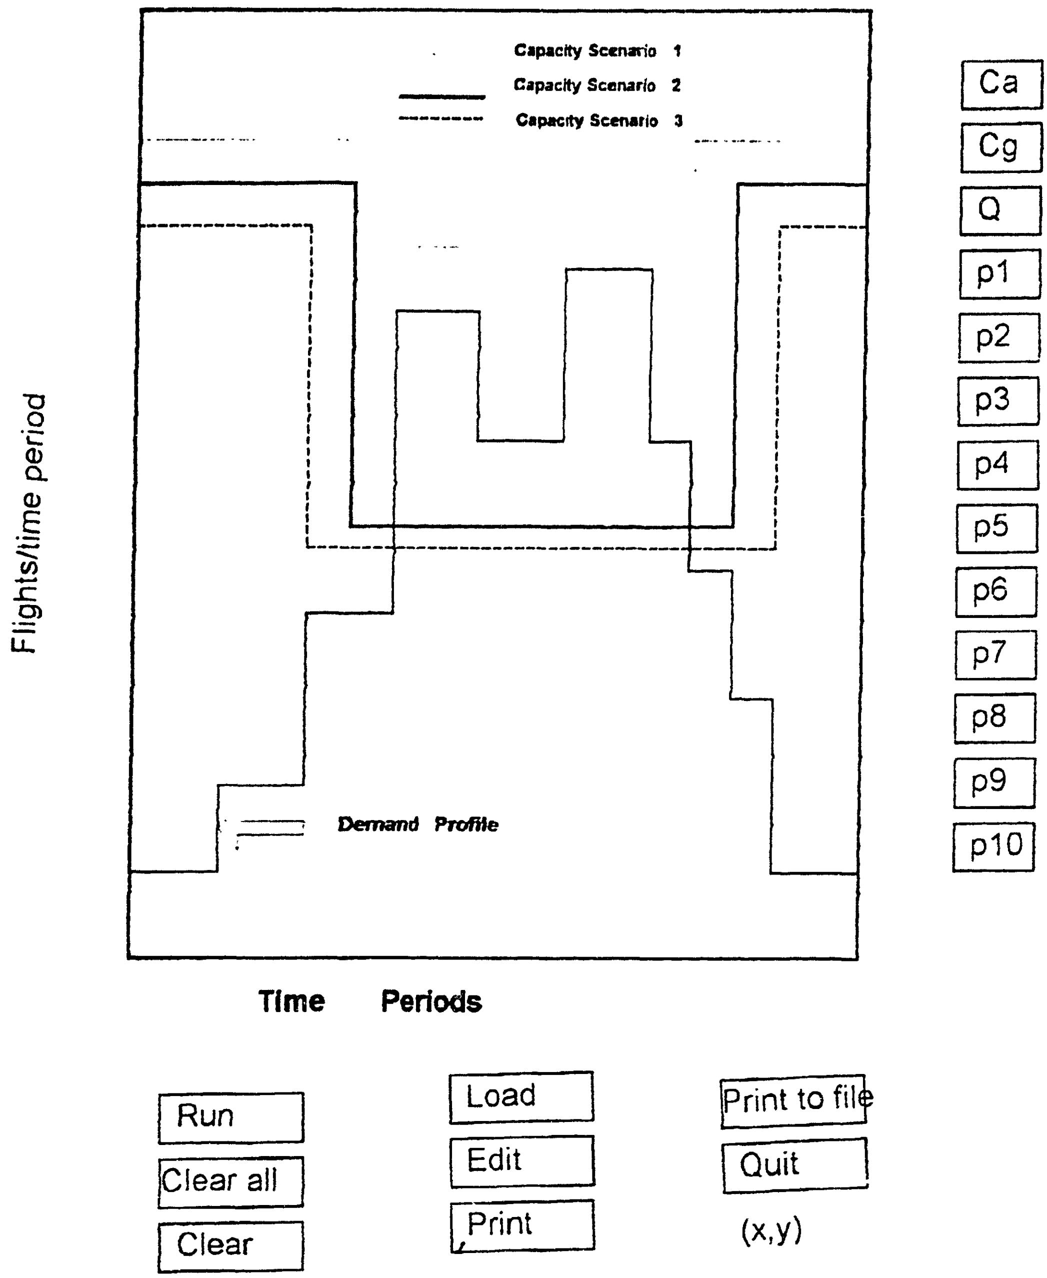

Fig. 3-6: Table of Delays and Costs in GUI

Air Delay
Ground Delay
Total Delay
Air Delay Cost
Ground Delay Cost
Total Delay Cost

Fig. 3-7: Graph of Number of Flights Delayed in Air
(Example where Q=3)

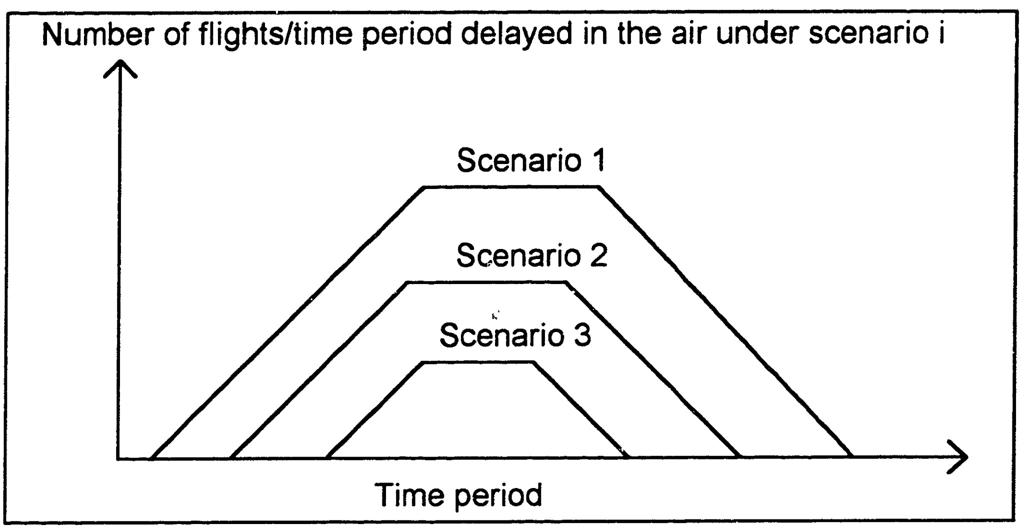

Chapter 4

Sensitivity Analysis and Results

The purpose of this chapter is to analyze the performance of the algorithm and to examine how it can be utilized in practice. This chapter is divided into five sections. The two cases of uncertainty in time and uncertainty in capacity are evaluated in each part. The first section takes a time period of capacity reduction and shifts it throughout the day. This is done in order to see how the effect of capacity reduction at an airport varies, depending on the time of day. We do this shift for three airports, each with different demands and levels of variability in their landing capacity. The second section investigates the sensitivity of the ground-holding algorithm to the number of capacity scenarios used, in terms of results and running time. We look at three different probability distributions for the capacity scenarios. In the third section, we analyze the sensitivity of the algorithm to T, the number of intervals the time horizon is divided into. We want to see the impact of T on the results and running time. The fourth section shows how the cost ratio of air delay to ground delay can influence the ground-holding policy. Finally in the fifth section, we compare our ground-holding algorithm to other scheduling stategies, in terms of amount of delay and cost.

4.1 Shift Period of Capacity Reduction Throughout the Day

In this section, we look at the two cases of capacity reduction, uncertainty in time and uncertainty in capacity, as they are applied to three airports. For each case, we will take a time period of capacity reduction and shift it throughout the day. This is done in order to see how the effect of capacity reduction on ground-holding is different at various times of the day. We also compare the ground-holding algorithm to the passive strategy of no ground-holding. One criterion we use for comparison, is the total cost of delays. We also examine the type of airport which might benefit the most from the algorithm. An airport that has a great deal of variability in its landing capacity may benefit more from the algorithm than an airport with a less volatile capacity. This may be because when there is less uncertainty in the capacity, the need for a good probabilistic model is not as great. This is the reason for selecting three different airports to find out what kind of demand and capacity profiles may benefit more from the algorithm.

The three airports we examine are Logan, LaGuardia, and San Francisco. We study the demand profile for arrivals at each of these airports from January 13, 1993. This appears to be a typical day for arrivals at each airport. These demands were taken from the Official Airline Guide(OAG) and are shown in Fig. 4-1a, 4-1b and 4-1c (see next page). The time horizon in these figures is 24 hours. The number of flights landing between midnight and 7 AM is small and therefore will be disregarded for the purpose of this study. This is based on the assumption that these flights will have no difficulty in landing. The time horizon for this study will then become 17 hours (7:01 AM to 12:00 AM(midnight)). We will divide this 17 hour horizon into 17 time windows, the first being 8 AM and the last one 12 AM. The 8 AM window includes all flights that are scheduled to land in the interval, 7:01 AM - 8:00 AM; 9 AM includes all flights scheduled between 8:01 AM - 9:00 AM; and so on. Each of these one hour time windows is further broken into four 15 minute time periods and T, the total number of time periods is calculated to be 68 (17 x 4 = 68). Time period 1 represents the time interval 7:01 - 7:15 AM, time period 2 represents 7:16 - 7:30 AM, and so on until time period 68, which represents 11:46 PM - 12:00 AM.

Fig. 4-1a: Profile of Daily Demand for Landing Operations at Logan Airport

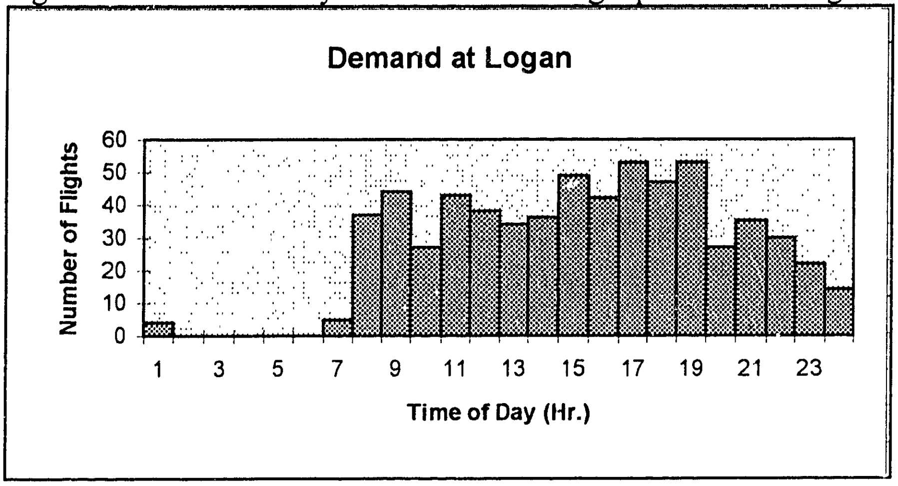

Fig. 4-1b: Profile of Daily Demand for Landing Operations at LaGuardia Airport

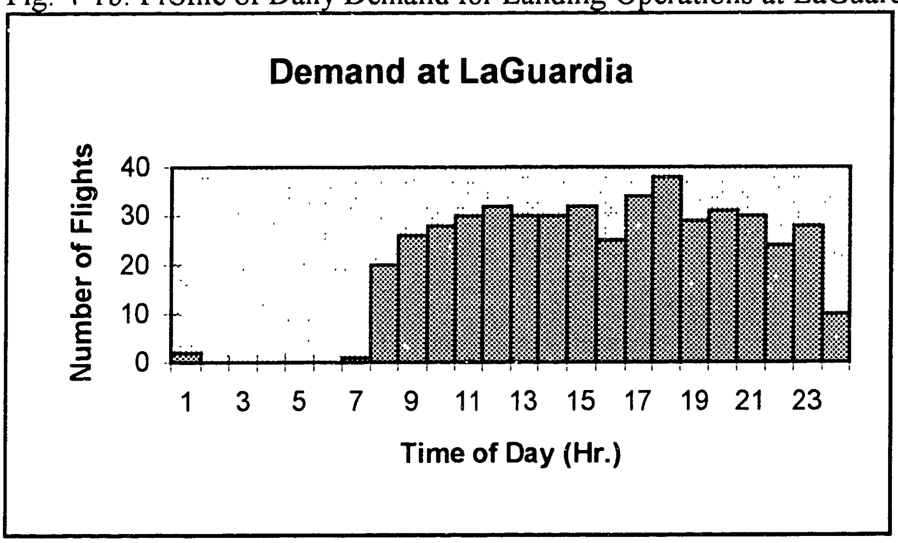

Fig. 4-1c: Profile of Daily Demand for Landing Operations at San Francisco Airport

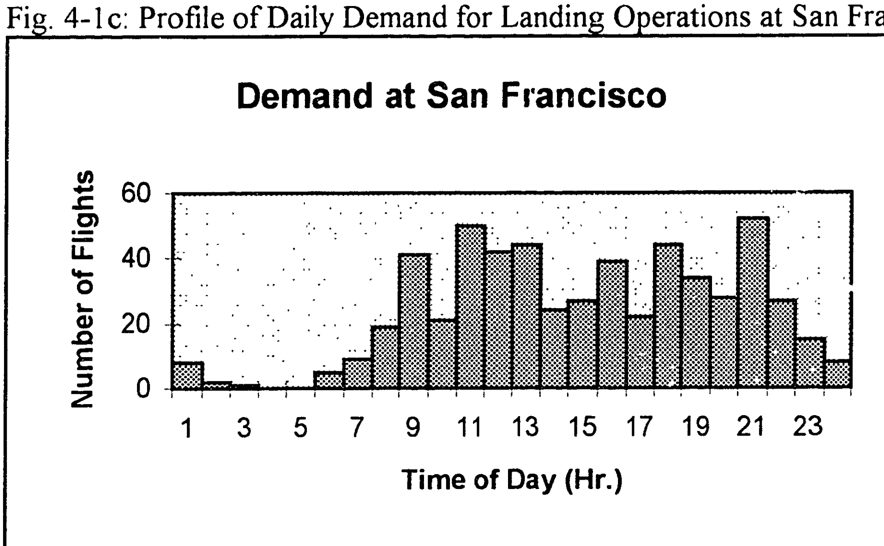

Depending on weather, landing procedures at an airport may be classified into four separate categories. These four categories are VFR1, VFR2, IFR1 and IFR2 and correspond to different levels of visibility and cloud ceiling. VFR and IFR stand for Visual and Instrument Flight Rules, respectively. When visibility and cloud ceiling are low, IFR conditions prevail, instrument approaches take place at the airport and landing intervals between aircraft are greater, typically making capacity lower than what it is under VFR. VFR and IFR are each split into two levels. VFR1 has the highest landing capacity, followed by VFR2, which is marginal VFR. Next is IFR1, followed by IFR2, which is low IFR.

Logan may operate under any of these four categories. At LaGuardia and San Francisco, three categories are identified: VFR, VFR/IFR, and IFR. The landing capacities at each airport are shown in Tables 4-1a, 4-1b and 4-1c. The number of arrivals that can land in an hour is converted to the number that can land in a 15 minute period and is rounded to an integer value. These values are based on the Engineered Performance Standards (EPS) and were obtained through the Draper Laboratory. The EPS numbers give capacity at airports according to Airport and Runway Configuration. It should be noted that these capacity values are approximate. This is not significant for the purpose of this experiment because we are interested in comparing airports with different levels of variability in their capacity profiles. The values used here serve this purpose.

Table 4-1a: Levels of Capacity at Logan

Weather Category	Arrivals/hr.	Arrivals/15 minutes
VFR1	60	15
VFR2	45	11
IFR1	36	9
IFR2	30	7

Table 4-1b: Levels of Capacity at LaGuardia

Weather Category	Arrivals/hr.	Arrivals/15 minutes
VFR	39	10
VFR/IFR	36	9
IFR	33	8

Table 4-1c: Levels of Capacity at San Francisco

Weather Category	Arrivals/hr.	Arrivals/15 minutes
VFR	53	13
VFR/IFR	41	10
IFR	26	6

We now define the cost of air and ground delays. If the cost of air delays is much greater than the cost of ground delays, then all delays will be ground-holding ones. If the cost of air delays is lower than that of ground delays, then all delays will be in the air. In practice, the cost of air delay is greater than that of ground delay. For the purpose of this study we use a ratio of air delay cost to ground delay cost of 5:3. This is a reasonably realistic ratio and it leads to strategies that do not have all delays in the air or all delays on the ground. To get an idea of the dollar value of these costs, we make the cost of air delays per time period, C_a, equal to \$1000/ hr. and the cost of ground-holding per period, C_g, equal to \$600/hr. Dividing the hour into four 15 minute periods, these costs become \$250/15 min. and \$150/15 min. periods. We use these costs for all examples in this study.

We shall now describe the example used in this section for case 1, uncertainty in time. Assume that a landing capacity reduction will occur at some time during the day and will last 2 hours. The capacity is at maximum during the rest of the day, except for this two-hour reduction, when capacity drops to its minimum value. For Logan, maximum is 15 arrivals/15 min. and minimum is 7 arrivals/15 min.; for LaGuardia, maximum is 10 arrivals/15 min. and minimum is 8 arrivals/15 min.; and for San Francisco, maximum is 13 arrivals/15 min. and minimum is 7 arrivals/15 min. This capacity reduction can begin at 3 possible times, each time equally likely. We start this example with the reduction occurring sometime at the beginning of the day and "shift" the reduction through the day, to finish with the reduction occurring at the end of the day. We obtain a total of 14 possible cases. The first shift means the reduction could start at period 1 (7:01-7:15 AM), 5 (8:01-8:15 AM), or 9 (9:01-9:15 AM) and would respectively, end at periods 8 (8:46-9:00 AM), 12 (9:46-10:00 AM), or 17 (10:46-11:00 AM). See Fig. 4-2.

Fig. 4-2: Shift 1 for Case 1, Uncertainty in Time

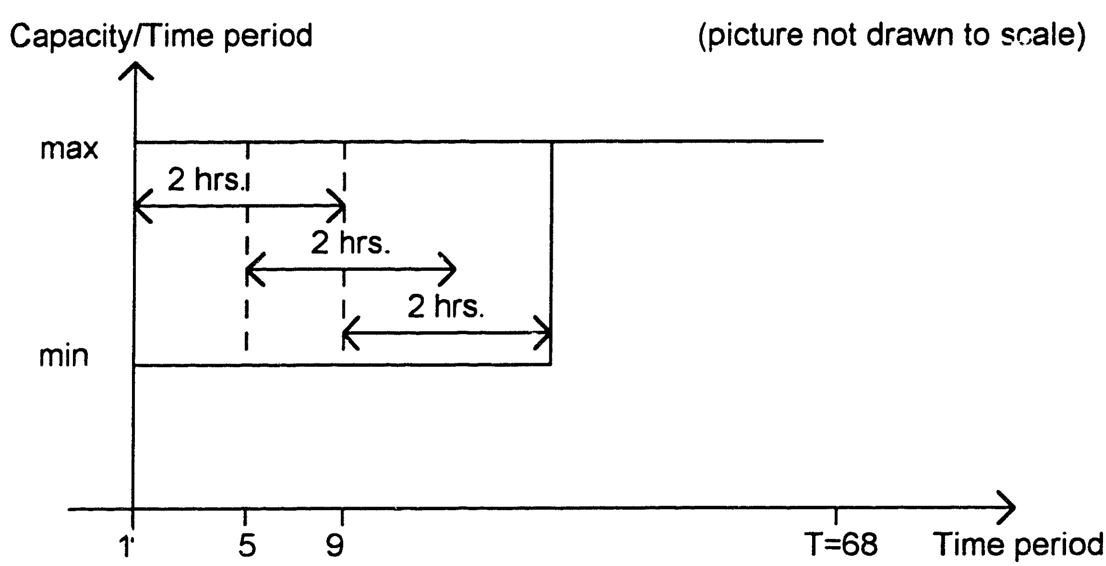

In Fig. 4-2, we can see that the reduction for shift 1 will last 2 hours and will begin at period 1, 5, or 9. In the second shift, the reduction could begin at period 5, 9, or 13 and will last 2 hours. See Fig. 4-3.

Fig. 4-3: Shift 2 for Case 1, Uncertainty in Time

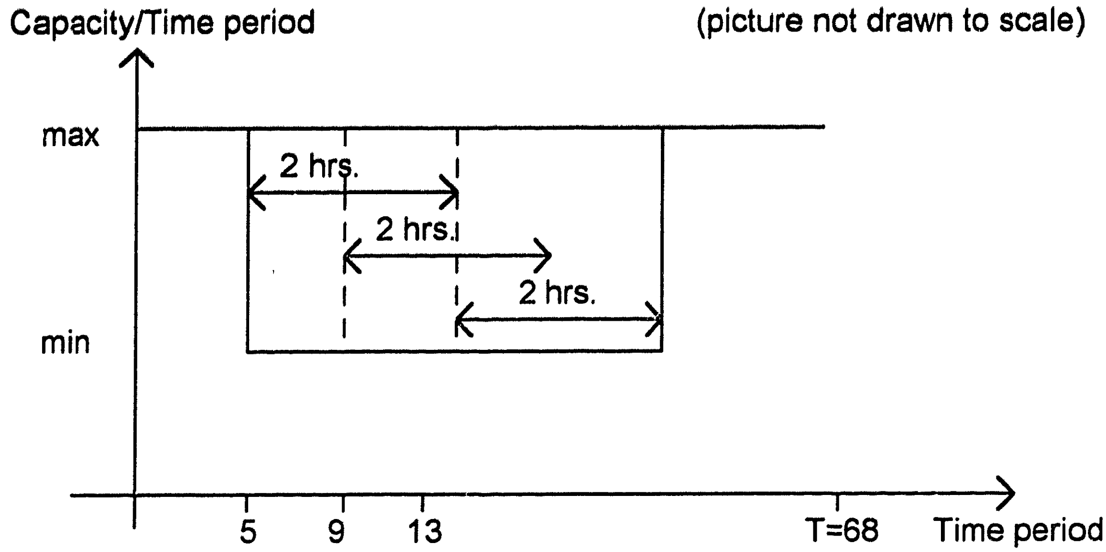

We then "shift" the reduction across the entire day to obtain 14 shifts. See Fig. 4-4.

Fig. 4-4: "Shift" Reduction Throughout Day for Case 1

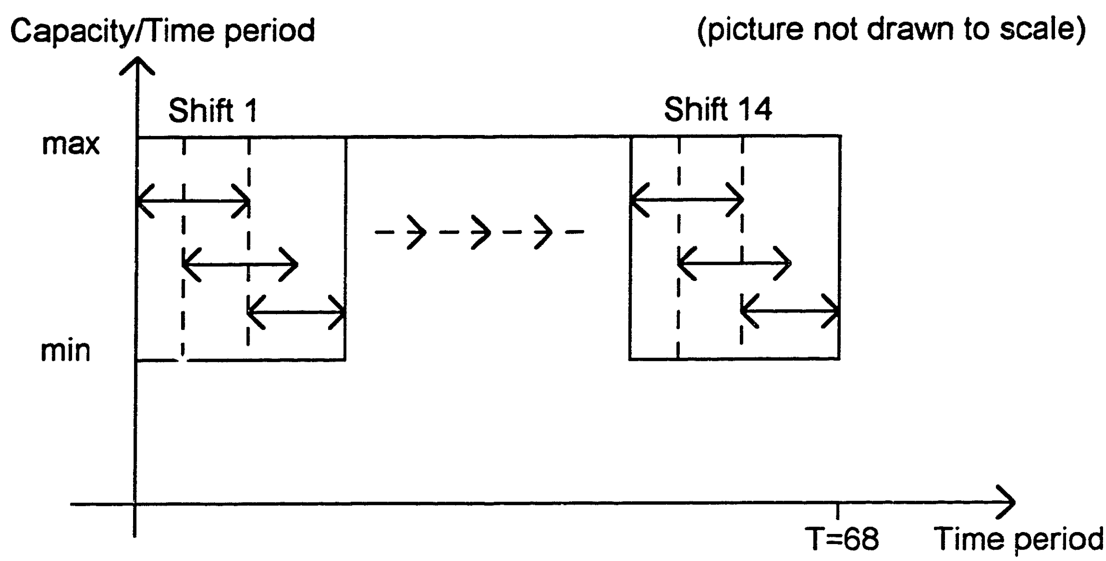

We do these 14 shifts for each of the three airports. The results can be seen in Table 4-2. The format of the table is as follows. There are 14 columns, each representing one of the 14 shifts. The header of the column indicates in which interval the reduction will occur. The top third of the table gives the results for Logan, the middle portion gives

LaGuardia's results, and the bottom third gives San Francisco's. All delays are given in hours and costs in dollars. For each airport, there are ten rows of values. The first three rows show the amount of expected air, ground, and total delay, respectively. The next three rows show their respective costs. The two rows after that give the amount of air delay, and associated cost, that would occur if a passive strategy, rather than the ground-holding algorithm, was used. A passive strategy means all flights take off according to schedule and all delays will be airborne delays. The last two rows show the ratio of the algorithm to the passive strategy, for total delay and total cost. Since the passive strategy gives the minimum amount of delay, we can see, percentage-wise, how much over the minimum amount of delay we are, when we use the algorithm. Even though the amount of total delay, using the algorithm, is more than the minimum possible, the cost is always less than or equal to that under a passive strategy.

Table 4-2: Results of "Shift" for Case 1, Uncertainty in Time

	7.01am-9.00am	8.01am-10.00am	9.01am-11.00am	10.01am-12.00p	11.01am-1.00pm	12.01pm-2.00pm	1.01pm-3.00pm	2.01pm-4.00pm	3.01pm-5.00pm	4.01pm-6.00pm	5.01pm-7.00pm	6.01pm-8.00pm	7.01pm-9.00pm	8.01pm-10.00pm
Logan Airport														
Air Delay (hrs.)	18.89	17.49	13.04	20.46	21.20	29.21	31.76	37.21	35.89	38.28	36.30	17.82	3.55	3.88
Ground Delay	7.25	12.00	16.00	9.00	13.50	39.25	66.75	71.75	62.25	39.00	5.50	4.00	4.00	1.50
Total Delay	26.14	29.49	29.04	29.46	34.70	68.46	98.51	108.96	98.14	77.28	41.80	21.82	7.55	5.38
Air Cost ($)	18892.5	17490.0	13035.0	20460.0	21202.5	29205.0	31762.5	37207.5	35887.5	38280.0	36300.0	17820.0	3547.5	3877.5
Ground Cost	4350.0	7200.0	9600.0	5400.0	8100.0	23550.0	40050.0	43050.0	37350.0	23400.0	3300.0	2400.0	2400.0	900.0
Total Cost	23242.5	24690.0	22635.0	25860.0	29302.5	52755.0	71812.5	80257.5	73237.5	61680.0	39600.0	20220.0	5947.5	4777.5
Passive:														
Air Delay (hrs.)	23.68	26.32	24.01	26.98	30.61	56.84	80.36	99.41	88.77	74.00	40.01	21.20	6.19	4.87
Air Cost ($)	23677.5	26317.5	24007.5	26977.5	30607.5	56842.5	80355.0	99412.5	88770.0	74002.5	40012.5	21202.5	6187.5	4867.5
Ground-holding Strategy/Passive Strategy:														
Total Delay	110%	112%	121%	109%	113%	120%	123%	110%	111%	104%	104%	103%	122%	110%
Total Cost	98%	94%	94%	96%	96%	93%	89%	81%	83%	83%	99%	95%	96%	98%
LaGuardia Airport	7.01am-9.00am	8.01am-10.00am	9.01am-11.00am	10.01am-12.00p	11.01am-1.00pm	12.01pm-2.00pm	1.01pm-3.00pm	2.01pm-4.00pm	3.01pm-5.00pm	4.01pm-6.00pm	5.01pm-7.00pm	6.01pm-8.00pm	7.01pm-9.00pm	8.01pm-10.00pm
Air Delay (hrs.)	0.00	0.00	0.00	0.00	0.00	0.00	0.33	2.23	1.65	2.89	2.31	0.00	0.00	0.00
Ground Delay	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	5.25	3.00	0.00	0.00	0.00	0.00
Total Delay	0.00	0.00	0.00	0.00	0.00	0.00	0.33	3.23	6.90	5.89	2.31	0.00	0.00	0.00
Air Cost ($)	0.0	0.0	0.0	0.0	0.0	0.0	330.0	2227.5	1650.0	2887.5	2310.0	0.0	0.0	0.0
Ground Cost	0.0	0.0	0.0	0.0	0.0	0.0	0.0	600.0	3150.0	1800.0	0.0	0.0	0.0	0.0
Total Cost	0.0	0.0	0.0	0.0	0.0	0.0	330.0	2827.5	4800.0	4687.5	2310.0	0.0	0.0	0.0
Passive:														
Air Delay (hrs.)	0.00	0.00	0.00	0.00	0.00	0.00	0.33	2.89	5.20	4.87	2.31	0.00	0.00	0.00
Air Cost ($)	0.0	0.0	0.0	0.0	0.0	0.0	330.0	2887.5	5197.5	4867.5	2310.0	0.0	0.0	0.0
Ground-holding Strategy/Passive Strategy														
Total Delay	100%	100%	100%	100%	100%	100%	100%	112%	133%	121%	100%	100%	100%	100%
Total Cost	100%	100%	100%	100%	100%	100%	100%	98%	92%	96%	100%	100%	100%	100%
San Francisco Airport	7.01am-9.00am	8.01am-10.00am	9.01am-11.00am	10.01am-12.00p	11.01am-1.00pm	12.01pm-2.00pm	1.01pm-3.00pm	2.01pm-4.00pm	3.01pm-5.00pm	4.01pm-6.00pm	5.01pm-7.00pm	6.01pm-8.00pm	7.01pm-9.00pm	8.01pm-10.00pm
Air Delay (hrs.)	19.31	29.95	25.74	38.61	31.43	17.24	7.51	15.02	20.54	20.38	27.97	15.43	29.45	18.15
Ground Delay	21.50	46.25	59.75	32.50	3.25	1.50	12.75	11.00	20.75	24.75	20.75	35.75	4.75	2.50
Total Delay	40.81	76.20	85.49	71.11	34.68	18.74	20.26	26.02	41.29	45.13	48.72	51.18	34.20	20.65
Air Cost ($)	19305.0	29947.5	25740.0	38610.0	31432.5	17242.5	7507.5	15015.0	20542.5	20377.5	27967.5	15427.5	29452.5	18150.0
Ground Cost	12900.0	27750.0	35850.0	19500.0	1950.0	900.0	7650.0	6600.0	12450.0	14850.0	12450.0	21450.0	2850.0	1500.0
Total Cost	32205.0	57697.5	61590.0	58110.0	33382.5	18142.5	15157.5	21615.0	32992.5	35227.5	40417.5	36877.5	32302.5	19650.0
Passive														
Air Delay (hrs.)	33.50	65.67	72.27	69.47	33.58	18.40	15.92	22.28	36.38	37.87	44.47	43.89	33.25	19.80
Air Cost ($)	33495.0	65670.0	72270.0	69465.0	33577.5	18397.5	15922.5	22275.0	36382.5	37867.5	44467.5	43890.0	33247.5	19800.0
Ground-holding Strategy/Passive Strategy.														
Total Delay	122%	116%	118%	102%	103%	102%	127%	117%	113%	119%	110%	117%	103%	104%
Total Cost	96%	88%	85%	84%	99%	99%	95%	97%	91%	93%	91%	84%	97%	99%

In Table 4-2, we can see for Logan, that the amount of delay and cost peaks in the afternoon. When we look at the ratio of the ground-holding algorithm to the passive strategy, we can see the total cost of ground-holding if the capacity reduction takes place in the afternoon is 10-20% less than that under a passive strategy. For LaGuardia, delays occur in only five shifts, those in the afternoon and early evening, and are quite small. Thus, the total cost under a ground-holding strategy is not much less than that under a passive strategy. (This also suggests that the capacity and/or demand levels we have used for LaGuardia in this experiment may be too high and low, respectively.) San Francisco appears to have a more erratic pattern of delays throughout the day. Most delays occur in the morning, then lessen noticeably in the early afternoon, increase again in the early evening, and taper off toward the end of the day. When using the ground-holding algorithm, higher savings in cost are associated with periods of larger amounts of delay.

In Fig. 4-5a, 4-5b, and 4-5c, we can see, for each of the airports, how much of the total delay is air delay and how much is ground delay and how this varies across the 14 shifts.

Fig. 4-5a: Case 1: Air and Ground Delay at Logan Airport

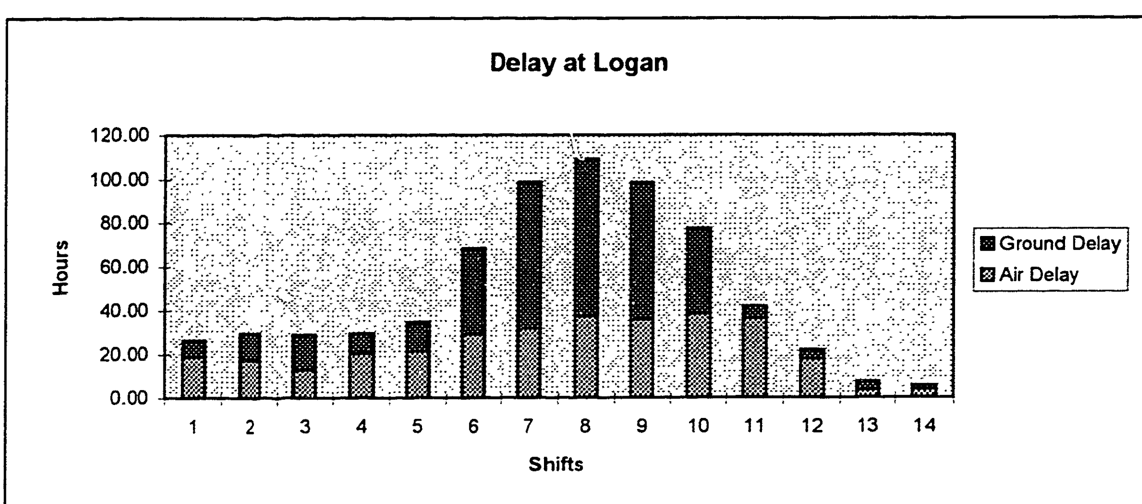

Fig. 4-5b: Case 1: Air and Ground Delay at LaGuardia Airport

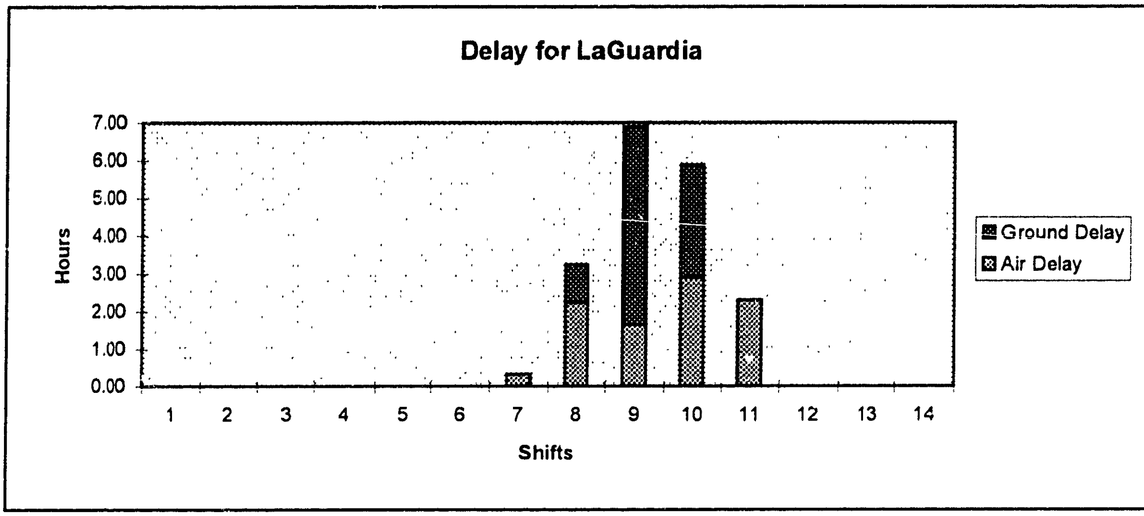

Fig. 4-5c: Case 1: Air and Ground Delay at San Francisco Airport

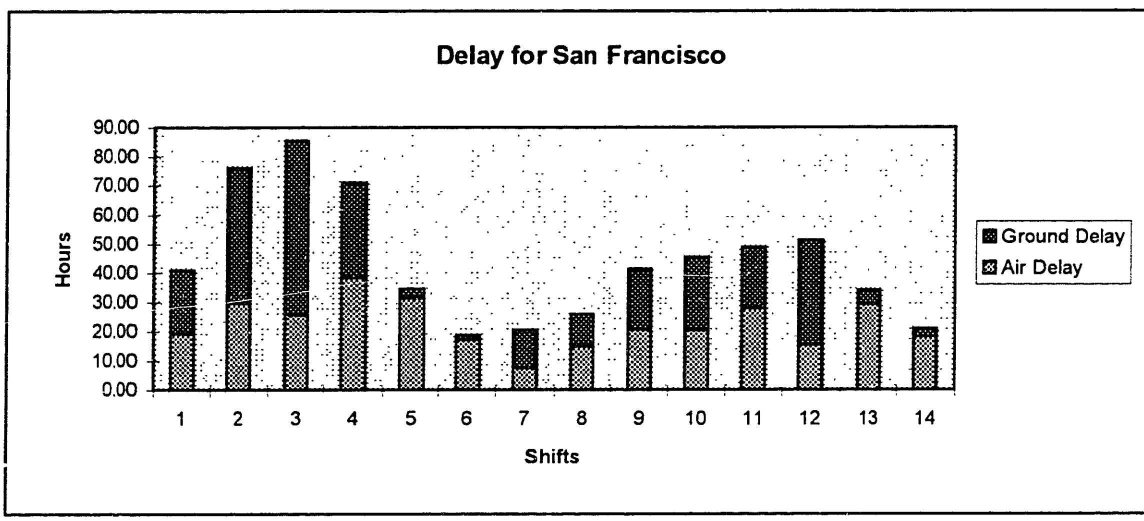

As seen in the above figures, there is a greater amount of ground delay than air delay whenever there is a large amount of total delay. This agrees with our intuition. When there is a small amount of delay, this delay will be air delay so that the system is not underused, meaning available landing opportunities are not wasted.

Combining the results of Table 4-2 and Figs. 4-5, we can conclude that the ground-holding strategy is always cheaper than the passive strategy. Time periods of high delay assign more of the delays as ground delays rather than air delays. And these times of high delay represent higher savings in cost when comparing the algorithm to that of the passive strategy.

We now describe the example used for case 2, uncertainty in capacity. The example is slightly different, though the concept of 14 shifts still exists. The capacity reduction will last four hours and the start (and end) time of this reduction is known with certainty, but the exact amount of reduction is unknown. Capacity outside this interval of reduction will be at maximum. For Logan, there are four possible capacity values during this interval, each equally probable. They are: 15 arrivals/15 min., 11 arrivals/15 min., 9 arrivals/15 min., or 7 arrivals/15 min. LaGuardia and San Francisco have three possible capacity values during the period of reduction, each equally likely. For LaGuardia, these values are 10 arrivals/15 min., 9 arrivals/15 min. or 8 arrivals/15 min. and for San Francisco, they are 13 arrivals/15 min., 10 arrivals/15 min., or 6 arrivals/15 min. As for case 1, we start the example with the reduction occurring in the beginning of the day and shift it throughout the day, to obtain a total of 14 shifts. This means the first shift will occur during periods 1-16 (7:01-11:00 AM) where capacity may be reduced. See Fig. 4-6 for an example of the first shift for San Francisco.

Fig. 4-6: Shift 1 for Case 2, Uncertainty in Capacity for San Francisco

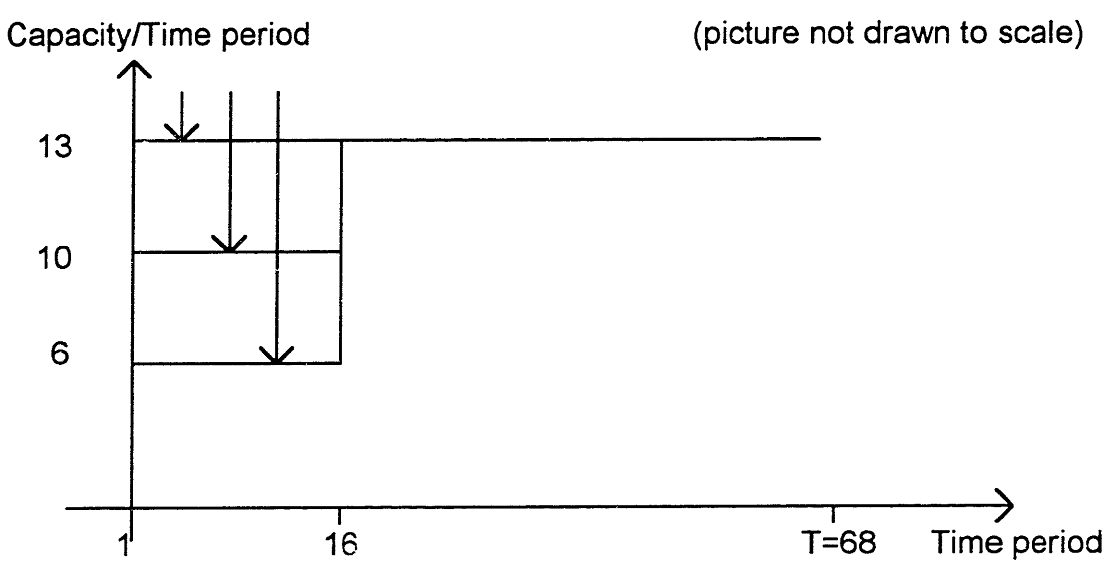

For San Francisco, under shift 1, capacity during periods 1-16 will be 13, 10, or 6. During the second shift, capacity will again be reduced, with equal likelihood to 13, 10, or 6 for periods 5-20. We continue this way through the day. See Fig. 4-7.

Fig. 4-7: Shift Reduction Through the Day for Case 2 for San Francisco

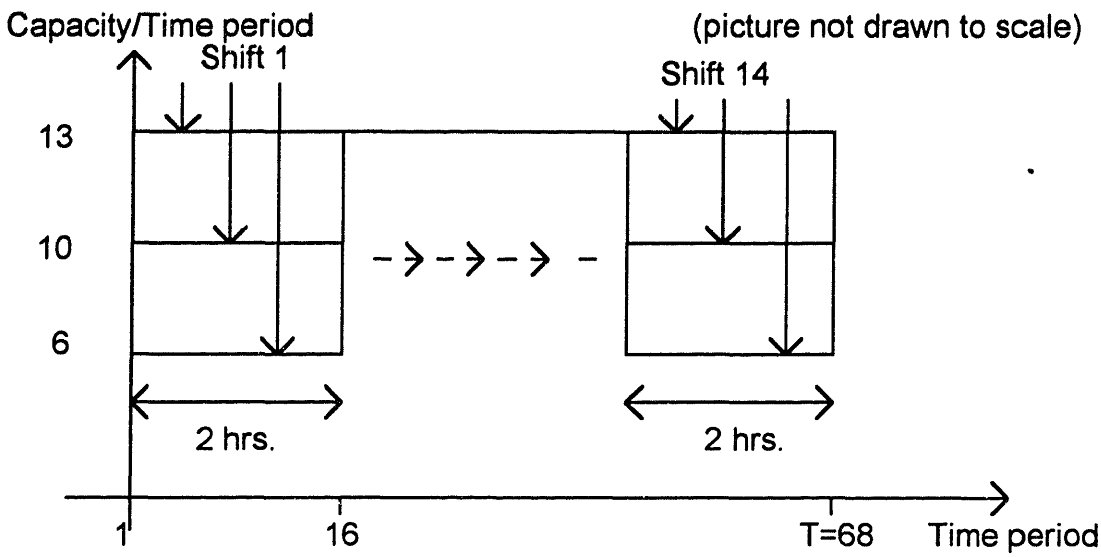

The results for case 2 can be seen in Table 4-3. It is set up in the same manner as Table 4-2, the table of results for case 1.

Table 4-3: Results of "Shift" for Case 2, Uncertainty in Capacity

	7:01am-11:00am	8:01am-12:00pm	9:01am-1:00pm	10:01am-2:00pm	11:01am-3:00pm	12:01pm-4:00pm	1:01pm-5:00pm	2:01pm-6:00pm	3:01pm-7:00pm	4:01pm-8:00pm	5:01pm-9:00pm	6:01pm-10:00pm	7:01pm-11:00pm	8:01pm-12:00am
Logan Airport														
Air Delay (hrs.)	30.69	32.50	23.63	43.63	38.00	55.06	81.31	98.69	79.63	75.69	54.44	32.63	4.63	4.63
Ground Delay	0.00	0.00	0.00	0.00	2.75	6.25	18.00	38.00	37.00	43.00	12.25	6.25	0.00	0.00
Total Delay	30.69	32.50	23.63	43.63	40.75	61.31	99.31	136.69	116.63	118.69	66.69	38.88	4.63	4.63
Air Cost ($)	30687.5	32500.0	23625.0	43625.0	38000.0	55062.5	81312.5	98687.5	79625.0	75687.5	54437.5	32625.0	4625.0	4625.0
Ground Cost	0.0	0.0	0.0	0.0	1650.0	3750.0	10800.0	22800.0	22200.0	25800.0	7350.0	3750.0	0.0	0.0
Total Cost	30687.5	32500.0	23625.0	43625.0	39650.0	58812.5	92112.5	121487.5	101825.0	101487.5	61787.5	36375.0	4625.0	4625.0
Passive:														
Air Delay (hrs.)	30.69	32.50	23.63	43.63	40.06	59.75	94.81	127.19	107.38	107.94	63.63	37.31	4.63	4.63
Air Cost ($)	30687.5	32500.0	23625.0	43625.0	40062.5	59750.0	94812.5	127187.5	107375.0	107937.5	63625.0	37312.5	4625.0	4625.0
Ground-holding Strategy/Passive Strategy														
Total Delay	100%	100%	100%	100%	102%	103%	105%	107%	109%	110%	105%	104%	100%	100%
Total Cost	100%	100%	100%	100%	99%	98%	97%	96%	95%	94%	97%	97%	100%	100%
LaGuardia Airport														
Air Delay (hrs.)	0.00	0.00	0.00	0.00	0.00	0.00	0.33	2.31	3.80	5.03	2.81	0.00	0.00	0.00
Ground Delay	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.75	0.75	0.75	0.75	0.00	0.00	0.00
Total Delay	0.00	0.00	0.00	0.00	0.00	0.00	0.33	3.06	4.55	5.78	3.56	0.00	0.00	0.00
Air Cost ($)	0.0	0.0	0.0	0.0	0.0	0.0	330.0	2310.0	3795.0	5032.5	2805.0	0.0	0.0	0.0
Ground Cost	0.0	0.0	0.0	0.0	0.0	0.0	0.0	450.0	450.0	450.0	450.0	0.0	0.0	0.0
Total Cost	0.0	0.0	0.0	0.0	0.0	0.0	330.0	2760.0	4245.0	5482.5	3255.0	0.0	0.0	0.0
Passive:														
Air Delay (hrs.)	0.00	0.00	0.00	0.00	0.00	0.00	0.33	2.81	4.29	5.53	3.30	0.00	0.00	0.00
Air Cost ($)	0.0	0.0	0.0	0.0	0.0	0.0	330.0	2805.0	4290.0	5527.5	3300.0	0.0	0.0	0.0
Ground-holding Strategy/Passive Strategy														
Total Delay	100%	100%	100%	100%	100%	100%	100%	109%	106%	105%	108%	100%	100%	100%
Total Cost	100%	100%	100%	100%	100%	100%	100%	98%	99%	99%	99%	100%	100%	100%
San Francisco Airport														
Air Delay (hrs.)	36.22	57.42	47.77	73.34	51.07	34.73	13.70	27.89	44.80	39.77	56.51	34.40	26.65	24.26
Ground Delay	9.25	23.75	33.75	37.00	5.75	2.50	0.00	2.50	3.00	3.00	12.00	12.00	12.00	12.00
Total Delay	45.47	81.17	81.52	110.34	56.82	37.23	13.70	30.39	47.80	42.77	68.51	46.40	38.65	36.26
Air Cost ($)	36217.5	57420.0	47767.5	73342.5	51067.5	34732.5	13695.0	27885.0	44797.5	39765.0	56512.5	34402.5	26647.5	24255.0
Ground Cost	5550.0	14250.0	20250.0	22200.0	3450.0	1500.0	0.0	1500.0	1800.0	1800.0	7200.0	7200.0	7200.0	7200.0
Total Cost	41767.5	71670.0	68017.5	95542.5	54517.5	36232.5	13695.0	29385.0	46597.5	41565.0	63712.5	41602.5	33847.5	31455.0
Passive:														
Air Delay (hrs.)	42.32	73.10	70.04	97.76	54.86	36.38	13.70	29.54	46.78	41.75	64.43	42.32	34.57	32.18
Air Cost ($)	42322.5	73095.0	70042.5	97762.5	54862.5	36382.5	13695.0	29535.0	46777.5	41745.0	64432.5	42322.5	34567.5	32175.0
Ground-holding Strategy/Passive Strategy														
Total Delay	107%	111%	116%	113%	104%	102%	100%	103%	102%	102%	106%	110%	112%	113%
Total Cost	99%	98%	97%	98%	99%	100%	100%	99%	100%	100%	99%	98%	98%	98%

The pattern of the amount of delay across the day is similar to case 1, but it appears that there is more air delay than ground delay assigned for case 2. The cost savings are therefore fewer. This can be visually seen in Figs. 4-8a, 4-8b, and 4-8c below.

Fig. 4-8a: Case 2: Air and Ground Delay at Logan Airport

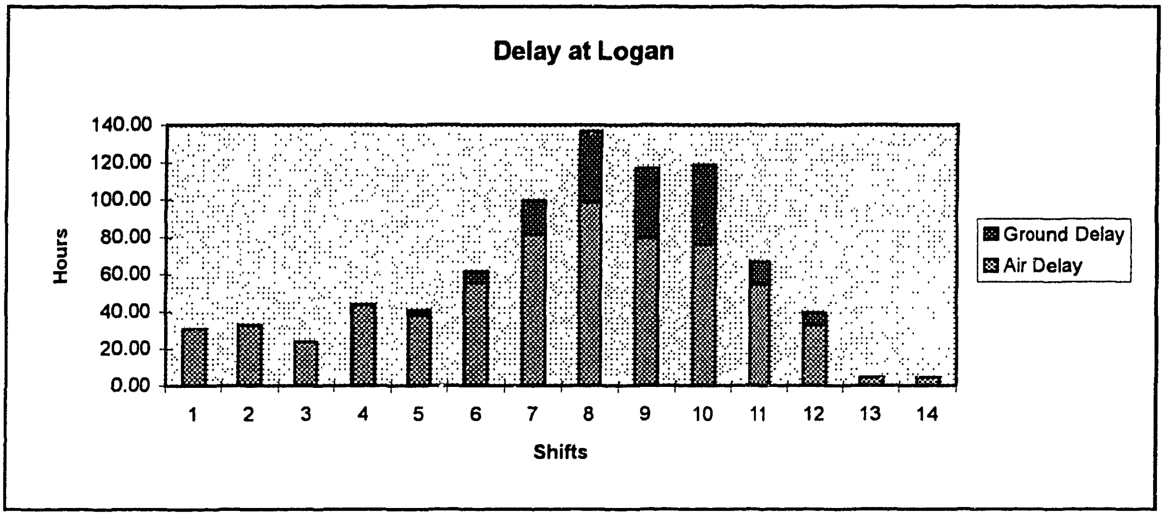

Fig. 4-8b: Case 2: Air and Ground Delay at LaGuardia Airport

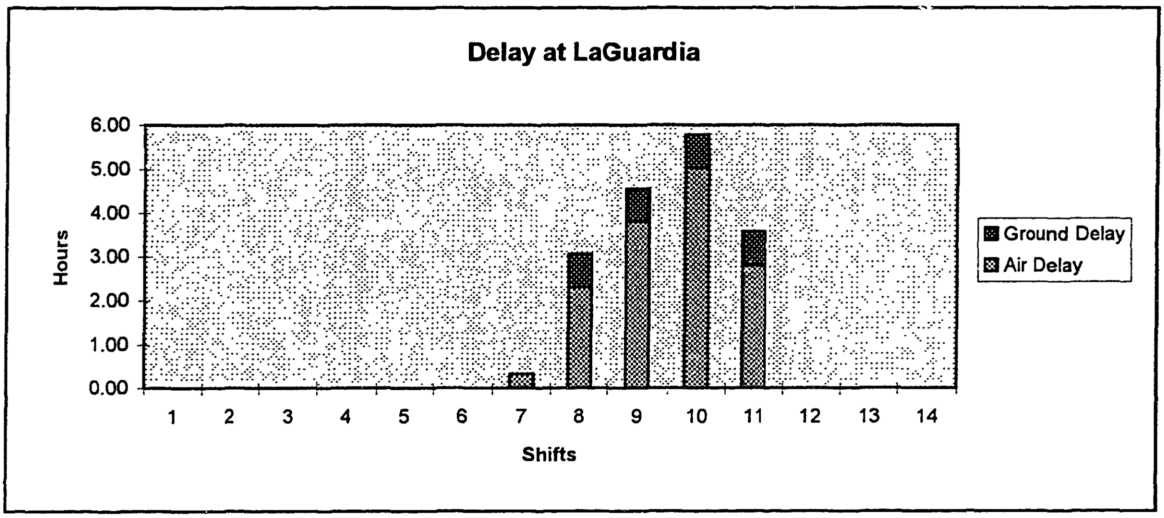

Fig. 4-8c: Case 2: Air and Ground Delay at San Francisco Airport

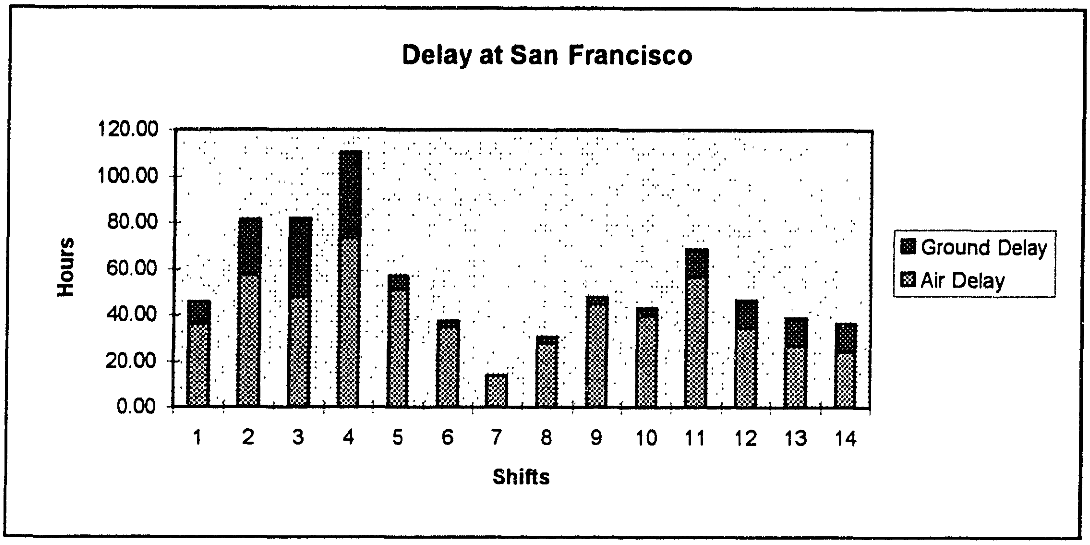

Table 4-3 and Figs. 4-8 seem to indicate that the cost savings of a ground-holding strategy for case 2 are less than that for case 1. One reason for this may be that for case 1, uncertainty in time, the capacity is definitely reduced to a minimum for two hours. Whereas for case 2, the capacity may or may not be reduced. The probability that capacity will be reduced to a minimum and the probability that capacity will stay at maximum is the same. The ground-holding algorithm will take this into account when assigning delays.

The examples done in this section show that when landing capacity does not vary much, in terms of maximum versus minimum capacity, ground-holding is not utilized as much. LaGuardia is an example of this. When there is a bigger discrepancy between the maximum and minimum capacity, then during times of high demand, there will be a great deal of ground-holding performed This will represent more cost-savings, in comparison to a passive strategy. Logan and San Francisco are examples of this. At these two airports, ground-holding was not necessary for the whole time horizon; it was only used in times of peak demand. We observe that at airports with situations where there is a high probability of demand exceeding available landing capacity, there is more ground-holding performed and this will represent higher savings.

4.2 Sensitivity to Number of Scenarios

The purpose of this section is to see how sensitive our algorithm is to Q, the number of probabilistic capacity scenarios used. We are interested in this analysis because Richetta used a Q equal to 3 or 4 in his tests. He believed a small value of Q was important for obtaining fast solutions and that it was consistent with current weather forecasting technology. Today, the kind of weather that will occur in a certain area can be predicted relatively accurately, but its severity and exact timing cannot. We wish to incorporate this uncertainty into a model and it is important to approximate the true probability distribution, associated with the timing and severity of a weather condition, as closely as possible. This means Q may have to be larger than 3 or 4. We wish to see how varying Q affects the results and solution time of the ground-holding problem.

First, we explain what we mean when we say Q can take on different values. In reality, time is continuous, but in our study, we divide the time horizon into time periods. If a weather condition, which will reduce landing capacity, is expected to arrive any time, with equal probability, in a 1 hour period, then in reality, there are infinite points of time in which this weather front may arrive. But if we divide the time horizon into 15 minute periods, then the weather front can arrive in only 4 possible time intervals. Each of these 4 intervals would represent a probabilistic capacity scenario and to represent the true probability distribution, Q would equal 4. Q is a discrete integer value because the time horizon is now comprised of time periods and the total number of time periods is an integer number. This situation would be an example of case 1, uncertainty in time.

For case 2, uncertainty in capacity, the time interval length will also affect how large Q can be, in order to accurately model the true probability distribution. For example, using the Logan capacity profile, the maximum number of flights that can land in an hour is 60 and the minimum is 30. If the time periods are 15 minutes long, then the capacity becomes the number of flights that can land in a 15 minute period. The maximum would then be 15 and the minimum around 7. If a weather forecast predicts capacity has an equally likely chance of being between maximum and minimum capacity for a certain duration, then (after conversion of the time horizon into 15 minute periods) there are 9 probabilistic capacity scenarios. The capacities between 7 and 15 each represent a scenario and Q equals 9. Again as in case 1, the number of possible capacity values is an integer because the possible capacity values are integer and finite.

In both cases, by discretizing the time horizon, we are changing the probability density function from continuous to discrete. When describing this probability distribution, Q should not be much larger than 10. If uncertainty lies in the arrival time of a weather front, current weather forecasting technology can usually pinpoint this time to within a couple of hours. Thus, if we use Logan as an example and use 15 minute time periods, then a capacity reduction that will begin sometime in a two hour period would mean a Q equal to 8. And as shown above for uncertainty in capacity, Q for Logan will not be larger than 9. This means, generally, Q should be less than or equal to 10. In this section, we want to see what happens when we try to approximate the probability distribution with fewer capacity scenarios, for problems involving Q equal to 9. And because solution time was a consideration in initially using a small Q, we want to see how much longer it takes to solve the problem for a larger Q.

The examples we do in this section use the Logan demand and capacity values that are described in the previous section. The time periods are 15 minutes long and the total number of periods is 68. For case 1, uncertainty in time, let us say that there will be a two hour storm that will reduce capacity from the maximum of 15 arrivals/15 min. to the minimum of 7 arrivals/15 min. In order to obtain a Q equal to 9, we will say that this storm will arrive sometime between 11:01 AM and 1:15 PM. Thus, the storm may arrive during one of the following nine periods: period 17 (11:01-11:15AM), period 18 (11:16-11:30AM), period 19 (11:31-11:45AM), period 20 (11:46AM-12:00PM), period 21 (12:01-12:15PM), period 22 (12:16-12:30PM), period 23 (12:31-12:45PM), period 24 (12:46-1:00PM), or period 25 (1:01-1:15PM). Each one of the above nine periods represents a probabilistic capacity scenario.

For case 2, uncertainty in capacity, we will say that there is a four hour storm which will begin at period 29 (2.01-2:15PM) and end at period 44 (5:46-6:00PM). During this time, capacity will be one of the following values (in arrivals/15 min.): 15, 14, 13, 12, 11, 10, 9, 8, or 7. When the storm occurs, each of the above values has a probability of being the capacity at the airport for that four hour interval. Thus, each possible capacity represents a probabilistic capacity scenario and like case 1, Q equals 9.

Using 15 minute periods, the maximum number of possible probabilistic capacity scenarios for both cases is 9. This means an accurate model of the true probability distribution would have 9 capacity scenarios, each with an associated probability. The first experiment we run in this section is to vary Q from 3 to 9, for both case 1 and 2, and then examine the results.

When Q is less than 9, only some of the original 9 scenarios are chosen. For example, in case 1, if we were to use Q equal to 3, then we might say the storm will arrive at period 17, 21, or 25. Table 4-4 shows what Q scenarios were chosen from the original nine when Q scenarios are used to solve the problem. This table is used for both case 1 and case 2. What we have done is label the capacity scenarios from 1 to 9. For case 1, scenario 1 refers to period 17, scenario 2 to period 18,..., and scenario 9 to period 25. For case 2, scenario 1 refers to capacity 15, scenario 2 to capacity 14,..., and scenario 9 to capacity 7. The left column shows what Q is and then each row shows what Q scenarios have been chosen to use in solving the problem. A dash in the table means that particular scenario was not chosen. For example, Q equal to 5 means we are using 5 probabilistic capacity scenarios in solving the problem and they are scenarios 1, 3, 5, 7, and 9. For case 1, this means choosing as the possible arrival time, periods 17, 19, 21, 23, and 25 and for case 2, the possible landing capacities during the storm would be 15, 13, 11, 9, and 7.

Table 4-4: Which Q Scenarios are Chosen for Each Q

Q	**Scenarios Chosen**								
9	1	2	3	4	5	6	7	8	9
8	1	2	3	4	-	6	7	8	9
7	1	2	-	4	5	6	-	8	9
6	1	2	-	4	-	6	-	8	9
5	1	-	3	-	5	-	7	-	9
4	1	-	-	4	-	6	-	-	9
3	1	-	-	-	5	-	-	-	9

For each case, we ran the example using three different probability distributions: uniform, triangular, and left-tail. For the uniform distribution, all scenarios had equal probability of occurring. If Q scenarios were used, then each scenario had a probability of 1/Q associated with it. For the triangular distribution, to calculate the probabilities, we used an isosceles triangle with base 1 and height 2 units. We divided the base into Q equal parts, each with length 1/Q. At every division, we vertically "sliced" the triangle to obtain Q slices and then found the area of each slice. The area of each of the Q slices was the probability associated with each of the Q scenarios. See Fig. 4-9.

Fig. 4-9: Calculation of Probabilities for Triangular Distribution

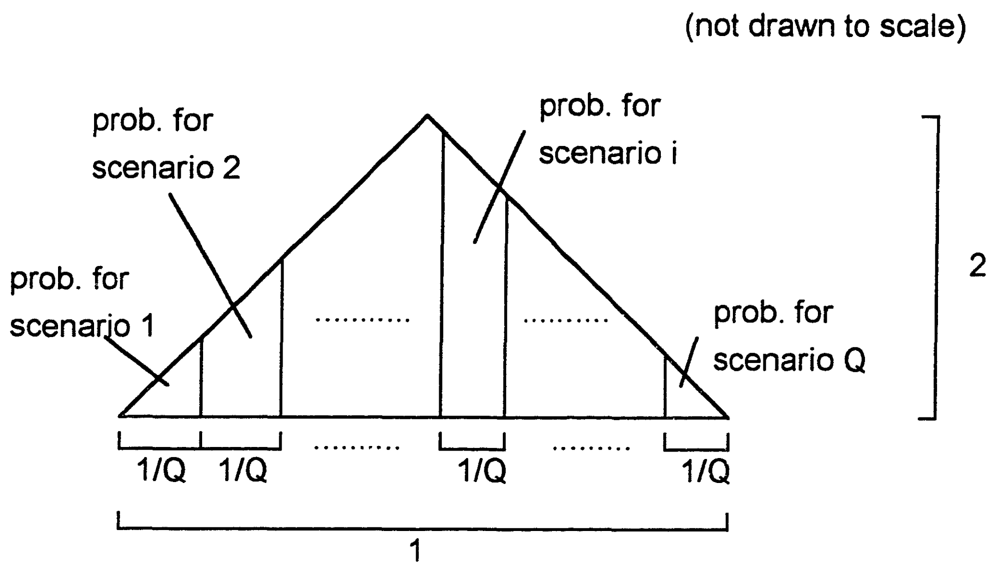

Table 4-5 shows the triangular probability distributions. The left column indicates the number of scenarios used. Each row shows which scenarios are used and their associated probabilities. P(i) is the probability of scenario i.

Table 4-5: Triangular Probability Distributions

Q	P(1)	P(2)	P(3)	P(4)	P(5)	P(6)	P(7)	P(8)	P(9)
9	2/81	6/81	10/81	14/81	17/81	14/81	10/81	6/81	2/81
8	1/32	3/32	5/32	7/32	0	7/32	5/32	3/32	1/32
7	2/49	6/49	0	10/49	13/49	10/49	0	6/49	2/49
6	1/18	3/18	0	5/18	0	5/18	0	3/18	1/18
5	2/25	0	6/25	0	9/25	0	6/25	0	2/25
4	1/8	0	0	3/8	0	3/8	0	0	1/8
3	2/9	0	0	0	5/9	0	0	0	2/9

For the left-tail distribution, we used a right angle triangle of base 1 and height 2 and calculated the probabilities of the scenarios in the same manner as the triangular distribution. Table 4-6 shows the probabilities of the left-tail distribution for the various values of Q. It is set up the same way as Table 4-5.

Table 4-6: Left-tail Probability Distributions

Q	P(1)	P(2)	P(3)	P(4)	P(5)	P(6)	P(7)	P(8)	P(9)
9	1/81	3/81	5/81	7/81	9/81	11/81	13/81	15/81	17/81
8	1/64	3/64	5/64	7/64	0	9/64	11/64	13/64	15/64
7	1/49	3/49	0	5/49	7/49	9/49	0	11/49	13/49
6	1/36	3/36	0	5/36	0	7/36	0	9/36	11/36
5	1/25	0	3/25	0	5/25	0	7/25	0	9/25
4	1/16	0	0	3/16	0	5/16	0	0	7/16
3	1/9	0	0	0	3/9	0	0	0	5/9

Figs. 4-10a, 4-10b, and 4-10c show, for case 1, the amount of air and ground delay for the various values of Q for the three probability distributions: uniform, triangular, and left-tail. Case 2 results are shown in Figs. 4-11a, 4-11b, and 4-11c. These figures show that ground and air delay assignments can vary significantly for the same example, depending on the number of probabilistic scenarios chosen to approximate the real distribution. Therefore, it is very important to select carefully the number of scenarios.

Fig. 4-10a: Case 1: Air and Ground Delay for Various Q for the Uniform Distribution

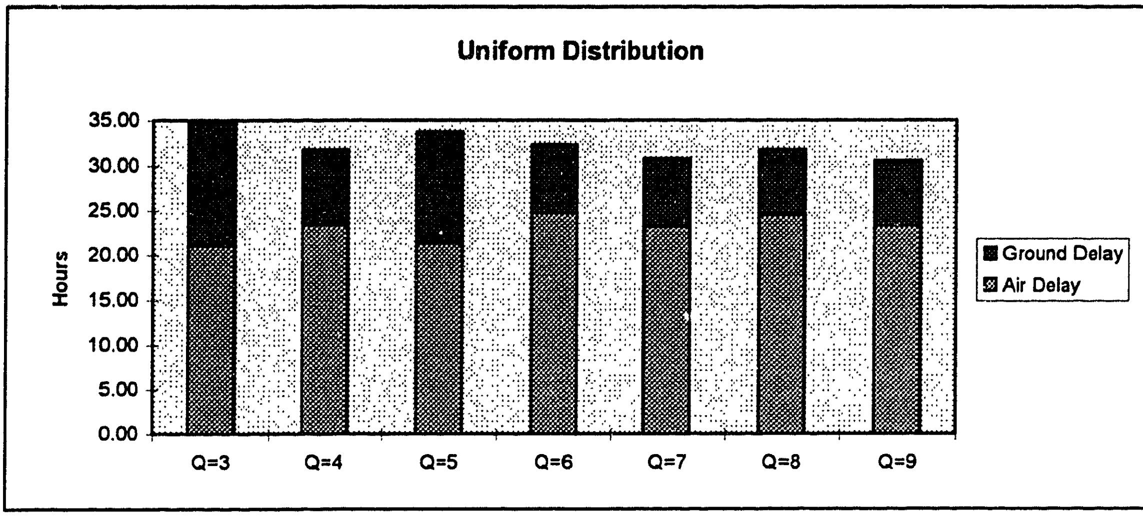

Fig. 4-10b: Case 1: Air and Ground Delay for Various Q for the Triangular Distribution

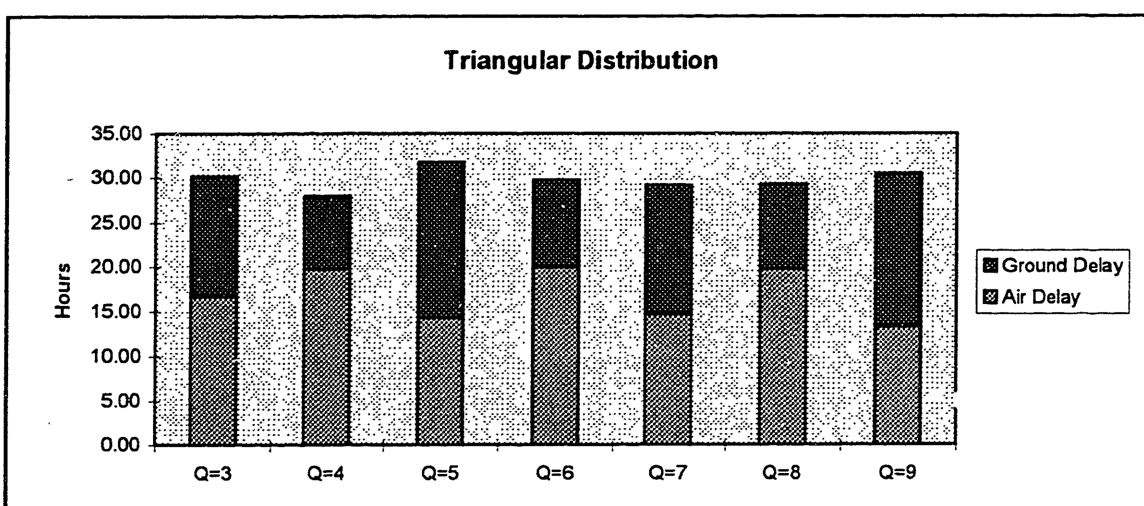

Fig. 4-10c: Case 1: Air and Ground Delay for Various Q for the Left-tail Distribution

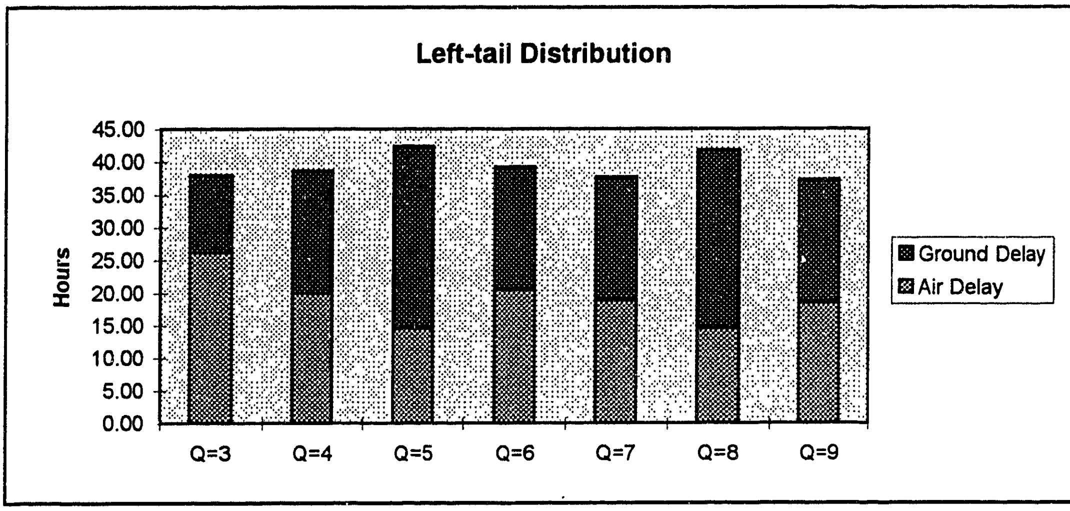

Fig. 4-11a: Case 2: Air and Ground Delay for Various Q for the Uniform Distribution

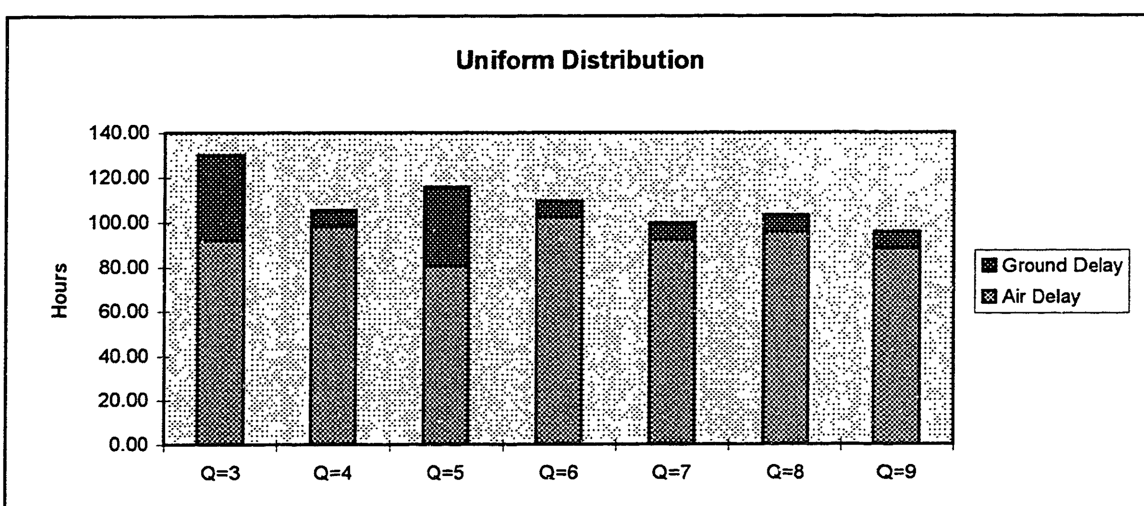

Fig. 4-11b: Case 2: Air and Ground Delay for Various Q for the Triangular Distribution

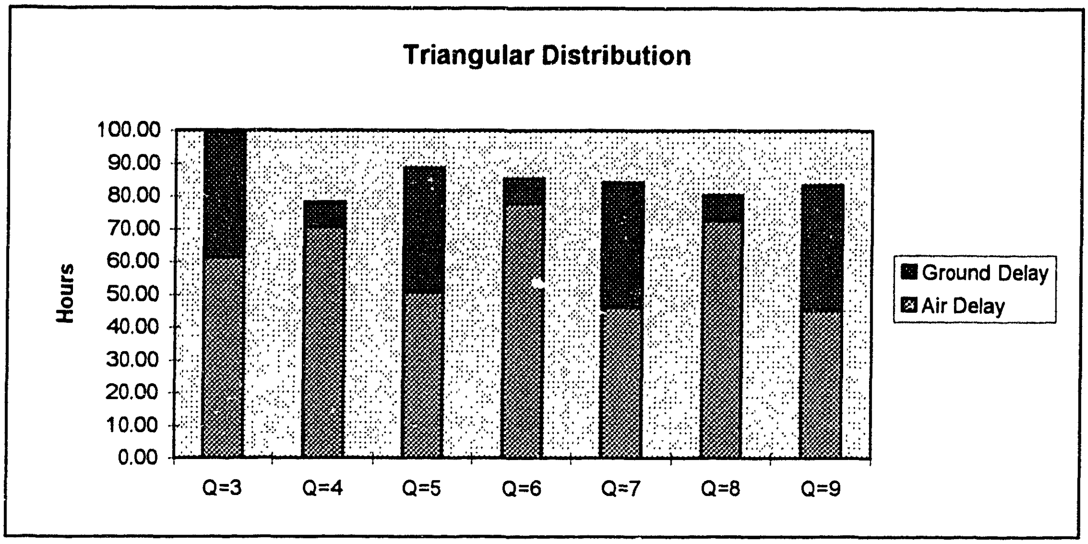

Fig. 4-11c: Case 2: Air and Ground Delay for Various Q for the Left-tail Distribution

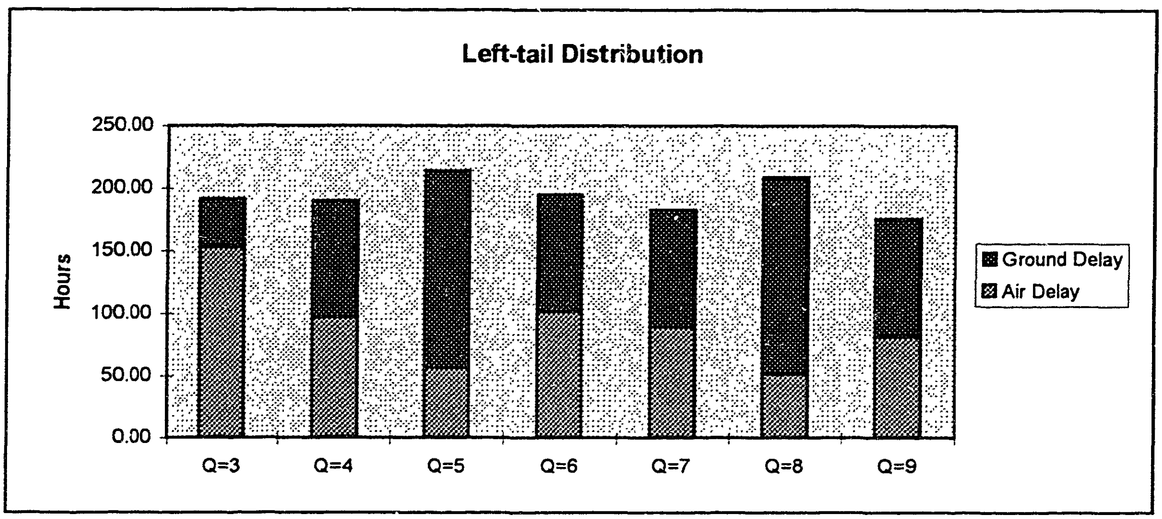

Another problem that arises when using less than 9 scenarios is how to choose which scenarios to use. If we use Q equal to 3, then there are four different logical ways of choosing 3 scenarios from the original 9. Each of these choices may result in different solutions. These four possible combinations are:

1.) Take the two outer scenarios and the middle one. (This is how we chose Q equal to 3 for the previous experiment in this section.)
   Scenarios chosen: 1, 5, 9
2.) Divide the nine scenarios into three intervals. Take the middle scenario of each interval.
   Scenarios chosen: 2, 5, 8
3.) Divide the nine scenarios into three intervals. Take the leftmost scenario of each interval.
   Scenarios chosen: 1, 4, 7
4.) Divide the nine scenarios into three intervals. Take the rightmost scenario of each interval.
   Scenarios chosen: 3, 6, 9

We will run this problem for case 1 and case 2, using each of the three probability distributions: uniform, triangular, and left-tail. The probability distribution associated with the three scenarios is the same as described for the previous example in this section. The results of these runs are shown in Fig. 4-12a, 4-12b, and 4-12c, for case 1, and in Figs. 4-13a, 4-13b, and 4-13c, for case 2. For each combination, we show the amount of air and ground delay. It can be seen that, depending on which scenarios are chosen, the amount of air and ground delay can vary.

Fig. 4-12a: Case 1: Air and Ground Delay for Q=3 Combinations for Uniform Distr.

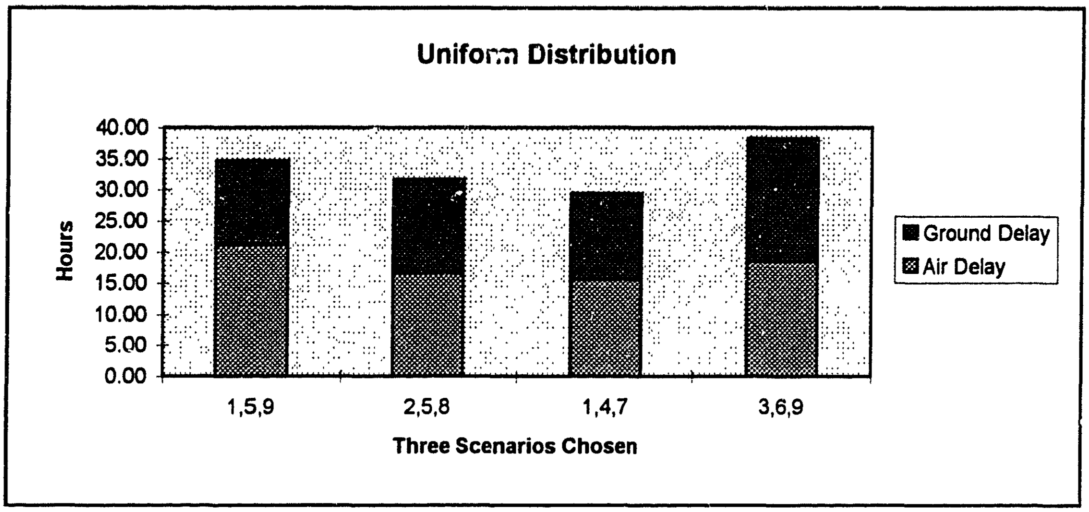

Fig. 4-12b: Case 1: Air and Ground Delay for Q=3 Combinations for Triangular Distr.

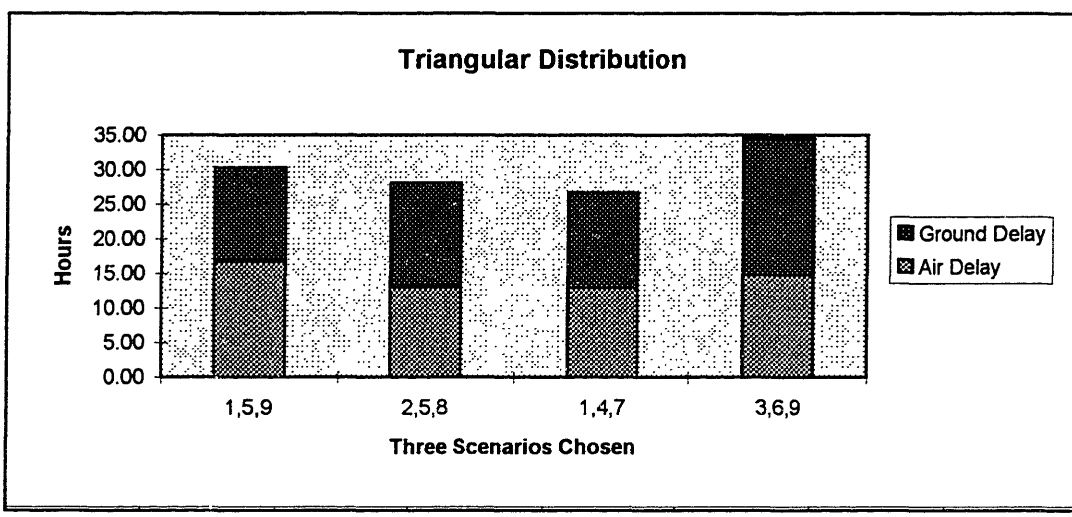

Fig. 4-12c: Case 1: Air and Ground Delay for Q=3 Combinations for Left-tail Distr.

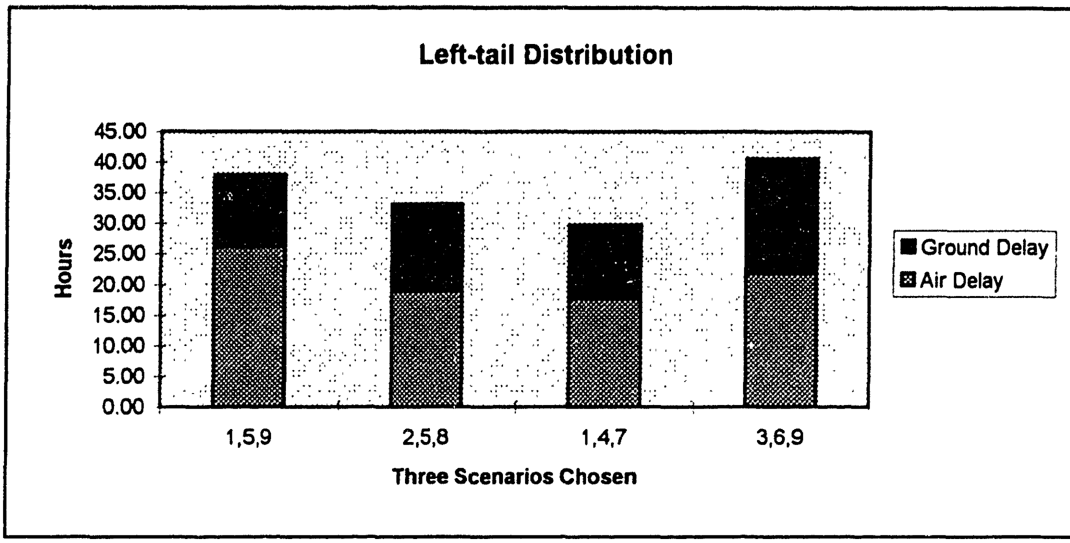

Fig. 4-13a: Case 2: Air and Ground Delay for Q=3 Combinations for Uniform Distr.

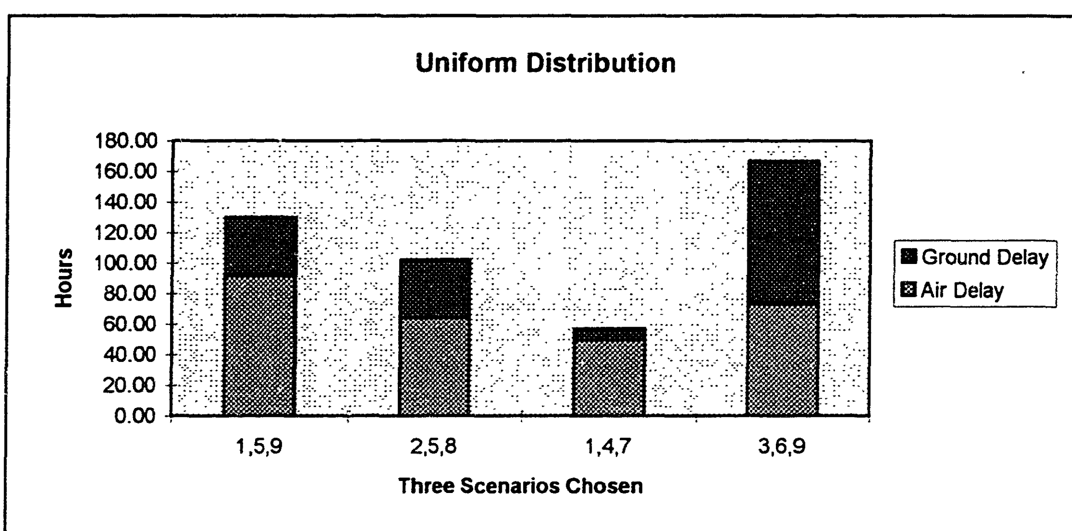

Fig. 4-13b: Case 2: Air and Ground Delay for Q=3 Combinations for Triangular Distr.

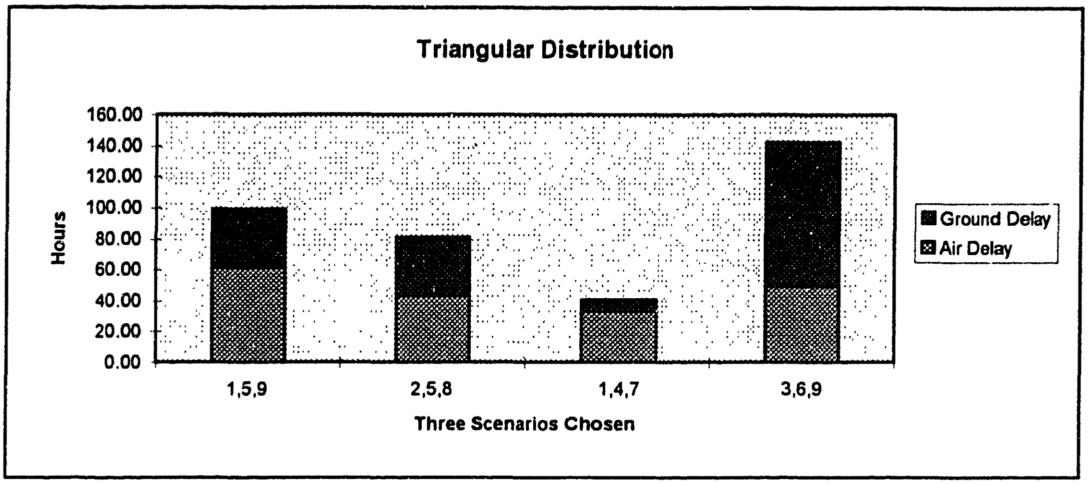

Fig. 4-13c: Case 2: Air and Ground Delay for Q=3 Combinations for Left-Tail Distr.

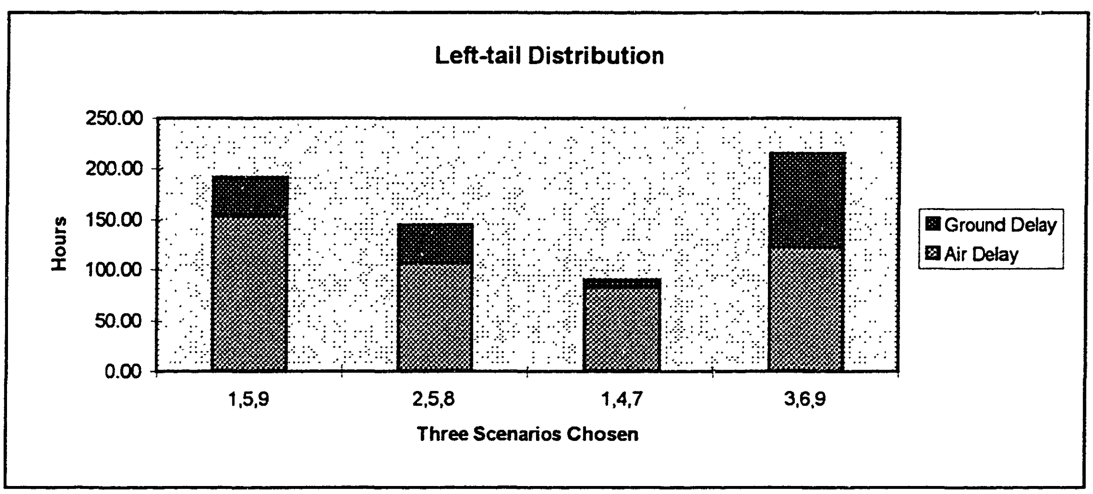

The above two experiments show how sensitive the ground-holding algorithm is to Q, the number of scenarios used. It is difficult to decide how many scenarios and which to use when we try to approximate the true distribution with fewer scenarios. It appears that it would be best to use the Q that is necessary to describe the true probability distribution. The question then is whether the time to solve the problem will be too long. Table 4-7 shows, for different values of Q, the time to generate the matrix, the time to solve the problem, and then the total time, which is equal to generation time plus solution time. Times are shown in seconds. Under each value of Q, the matrix size, [number of constraints] x [number of variables] is shown. Though we showed no examples with Q equal to 10, we ran some in order to get computation times for comparison.

Table 4-7: Generation, Solution, Total Times for Various Values of Q

	Q=3	Q=4	Q=5	Q=6	Q=7	Q=8	Q=9	Q=10
Mat Size	278x2825	348x2962	418x3099	488x3236	558x3373	628x3510	698x3647	768x3784
Generate	0.7	0.9	1.2	1.4	1.6	1.9	2.2	2.5
Solve	3.0	3.9	5.1	6.3	7.4	8.7	10.0	11.8
Total	3.7	4.8	6.3	7.7	9.0	10.6	12.2	14.3

Table 4-7 shows times for a T equal to 68, which represents a 17 hour day divided into 15 minute periods. For Q equal to 3, the computation time is about 4 seconds and for Q equal to 10, the time is about 14 seconds. Using Q equal to 10, more than triples the time, compared to Q equal to 3; but the time is still relatively short. As we mentioned in the beginning of this section, Q should, in practice, not get much larger than 10. Therefore, if 10 is the number of scenarios needed to describe the probability distribution, then that is how many scenarios should be used to obtain accurate results. In terms of the computation time required, this is a realistic and practical approach.

4.3 Sensitivity to Number of Time Intervals

The purpose of this section is to examine the sensitivity of the algorithm to T, the total number of time periods that the time horizon is divided into. This will help us in determining how long the length of each time period should be (i.e., 10 minute periods vs. 30 minute periods) in order to get a good solution. The smaller the time period is, the larger T will be. We want to see how the time to solve the problem increases as T increases. The computation time is dependent on the matrix size, [number of constraints] x [number of variables]. As we saw in Chapter 3, the matrix size is more dependent on T than Q. So even though we observed in the previous section that the computation time is still relatively small in relation to an increase in Q, this may not hold true for T. A large T may represent a large solution time.

The example solved in this section is similar to the one in the first section of this chapter, except we only look at Logan airport and Q equals 3 for both case 1 and case 2. We take an interval of capacity reduction and shift it across the day, which is 17 hours long, to obtain 14 shifts. For case 1, uncertainty in time, there will be a two-hour capacity reduction where capacity will drop from 60 arrivals/hr. to 30 arrivals/hr. This reduction could begin at 3 possible times, each an hour apart, and these different start times represent the 3 capacity scenarios. For case 2, uncertainty in capacity, there is a four-hour capacity reduction, where capacity for these four hours could be 60 arrivals/hr., 36 arrivals/hr., or 30 arrivals/hr. These 3 capacities each represent the three probabilistic capacity scenarios.

For case 1 and case 2, we do this example of 14 shifts using four different values of T, and then compare the amount of air and ground delay for each T. The four values of T represent time periods of length 10 minutes, 15 minutes, 20 minutes, and 30 minutes, and are, respectively, 102 (T=17*6), 68 (T=17*4), 51 (T=17*3), and 34 (T=17*2). Because Q equals 3 for both cases, we use only three weather categories for determining landing capacity. Since we are comparing across different time period lengths, the capacity must be consistent for the different values of T. Table 4-8 shows the weather category and corresponding capacity used for each time period length. The number of arrivals that can land in an hour are converted into how many can land in each time period.

Table 4-8: Capacity Under Various Time Period Lengths

Weather Category	# Arrivals/hour	# Arrivals/30 minutes	# Arrivals/20 minutes	# Arrivals/15 minutes	# Arrivals/10 minutes
VFR1	60	30	20	15	10
IFR1	36	18	12	9	6
IFR2	30	15	10	7.5	5

The number of arrivals per 15 minutes under weather category IFR2 is 7.5. Because all capacity values are restricted to be integer and we want to be consistent, we modeled 7.5 in the following manner. We alternated 7 and 8 for as many periods as the capacity was at 7.5. For example, if capacity was 7.5 for 4 time periods, then the capacity was set to 7, 8, 7, and 8 for the 4 periods.

The results from this example for case 1 and case 2 can be seen, respectively, in Table 4-9 and Table 4-10. The format of the tables is as follows. The first row indicates which shift or part of the day the reduction occurs in. Under each column, the air, ground, and total delay (in hrs.) is shown for the 10, 15, 20, and 30 minute periods. For each shift, we can see that the delay times for the different time period lengths are very similar. There is not much discrepancy in the delay times across the various values of T.

Table 4-9: Case 1: Delays for Various Values of T

Logan Airport	7:01am-9:00am	8:01am-10:00am	9:01am-11:00am	10:01am-12:00p	11:01am-1:00pm	12:01pm-2:00pm	1:01pm-3:00pm
10 min. periods:							
Air Delay (in hrs.)	13.59	13.48	10.56	15.62	16.23	24.15	27.61
Ground Delay (in hrs.)	6.33	9.67	11.50	5.33	7.83	29.83	54.50
Total Delay (in hrs.)	19.92	23.15	22.06	20.95	24.06	53.98	82.11
15 min. periods:							
Air Delay	14.77	14.19	10.73	15.76	16.50	24.59	27.97
Ground Delay	6.25	10.00	12.25	6.25	8.25	29.75	54.50
Total Delay	21.02	24.19	22.98	22.01	24.75	54.34	82.47
20 min. periods:							
Air Delay	14.63	14.08	11.00	15.84	16.83	24.86	27.94
Ground Delay	6.00	10.00	12.00	6.00	7.67	29.33	55.00
Total Delay	20.63	24.08	23.00	21.84	24.50	54.19	82.94
30 min. periods:							
Air Delay	14.85	14.52	10.89	17.16	17.49	26.07	28.71
Ground Delay	6.50	10.50	13.00	5.00	7.50	28.50	54.50
Total Delay	21.35	25.02	23.89	22.16	24.99	54.57	83.21
	2:01pm-4:00pm	3:01pm-5:00pm	4:01pm-6:00pm	5:01pm-7:00pm	6:01pm-8:00pm	7:01pm-9:00pm	8:01pm-10:00pm
10 min. periods:							
Air Delay (in hrs.)	32.89	32.01	33.72	29.54	14.36	2.20	2.48
Ground Delay (in hrs.)	61.50	53.17	33.00	6.17	2.50	1.67	0.00
Total Delay (in hrs.)	94.39	85.18	66.72	35.71	16.86	3.87	2.48
15 min. periods:							
Air Delay	33.25	31.93	33.99	29.95	14.93	2.48	2.64
Ground Delay	61.50	54.25	33.75	6.75	3.00	2.50	0.75
Total Delay	94.75	86.18	67.74	36.70	17.93	4.98	3.39
20 min. periods:							
Air Delay	33.22	32.56	34.10	29.70	14.63	2.75	2.64
Ground Delay	62.00	53.67	33.67	7.33	3.00	1.33	0.00
Total Delay	95.22	86.23	67.77	37.03	17.63	4.08	2.64
30 min. periods:							
Air Delay	33.66	33.66	34.49	31.02	14.85	3.30	2.81
Ground Delay	61.50	52.50	34.00	6.00	3.50	1.00	0.00
Total Delay	95.16	86.16	68.49	37.02	18.35	4.30	2.81

Table 4-10: Case 2: Delays for Various Values of T

Logan Airport	7:01am-11:00am	8:01am-12:00pm	9:01am-1:00pm	10:01am-2:00pm	11:01am-3:00pm	12:01pm-4:00pm	1:01pm-5:00pm
10 min. periods:							
Air Delay (in hrs.)	22.33	21.01	12.87	27.12	30.20	34.38	38.28
Ground Delay (in hrs.)	12.33	18.00	18.83	26.83	11.50	42.33	101.00
Total Delay (in hrs.)	34.66	39.01	31.70	53.95	41.70	76.71	139.28
15 min. periods:							
Air Delay	22.77	21.29	13.12	27.31	30.53	34.82	38.69
Ground Delay	13.50	19.50	19.25	27.25	11.50	42.00	100.75
Total Delay	36.27	40.79	32.37	54.56	42.03	76.82	139.44
20 min. periods:							
Air Delay	22.44	21.12	12.87	27.17	30.25	34.65	38.39
Ground Delay	13.67	19.67	19.67	27.67	12.33	42.67	102.00
Total Delay	36.11	40.79	32.54	54.84	42.58	77.32	140.39
30 min. periods:							
Air Delay	22.61	21.12	12.71	27.23	30.20	34.82	38.61
Ground Delay	14.50	21.00	21.00	28.50	13.50	43.00	102.00
Total Delay	37.11	42.12	33.71	55.73	43.70	77.82	140.61

	2:01pm-6:00pm	3:01pm-7:00pm	4:01pm-8:00pm	5:01pm-9:00pm	6:01pm-10:00pm	7:01pm-11:00pm	8:01pm-12:00am
10 min. periods:							
Air Delay (in hrs.)	37.57	32.95	29.81	23.60	18.70	3.03	3.03
Ground Delay (in hrs.)	157.00	127.17	131.17	70.50	31.00	0.00	0.00
Total Delay (in hrs.)	194.57	160.12	160.98	94.10	49.70	3.03	3.03
15 min. periods:							
Air Delay	37.95	33.41	30.28	24.01	19.06	3.38	3.38
Ground Delay	157.00	127.25	132.00	71.50	32.00	0.00	0.00
Total Delay	194.95	160.66	162.28	95.51	51.06	3.38	3.38
20 min. periods:							
Air Delay	37.62	33.00	29.92	23.54	18.48	2.86	2.86
Ground Delay	158.33	128.33	132.33	71.67	32.00	0.00	0.00
Total Delay	195.95	161.33	162.25	95.21	50.48	2.86	2.86
30 min. periods:							
Air Delay	37.62	32.84	29.87	23.76	18.81	2.97	2.97
Ground Delay	159.00	129.00	133.00	72.00	32.00	0.00	0.00
Total Delay	196.62	161.84	162.87	95.76	50.81	2.97	2.97

In order to visualize what is happening in Tables 4-9 and 4-10, we plot the air and ground delays for cases 1 and 2. Figs. 4-14a and 4-14b show, respectively, the air and ground delay for case 1, and Figs. 4-15a and 4-15b show the air and ground delay for case 2. For each shift, the delay for the 4 different time period lengths is charted. In each shift, the first vertical bar corresponds to the 10 minute time periods, the second to the 15 minute periods, the third to the 20 minute periods, and the fourth to the 30 minute periods. These graphs show that for both case 1 and case 2, the air and ground delays for the different values of T differ by only a small amount.

Fig. 4-14a: Case 1: Air Delay for Various Values of T

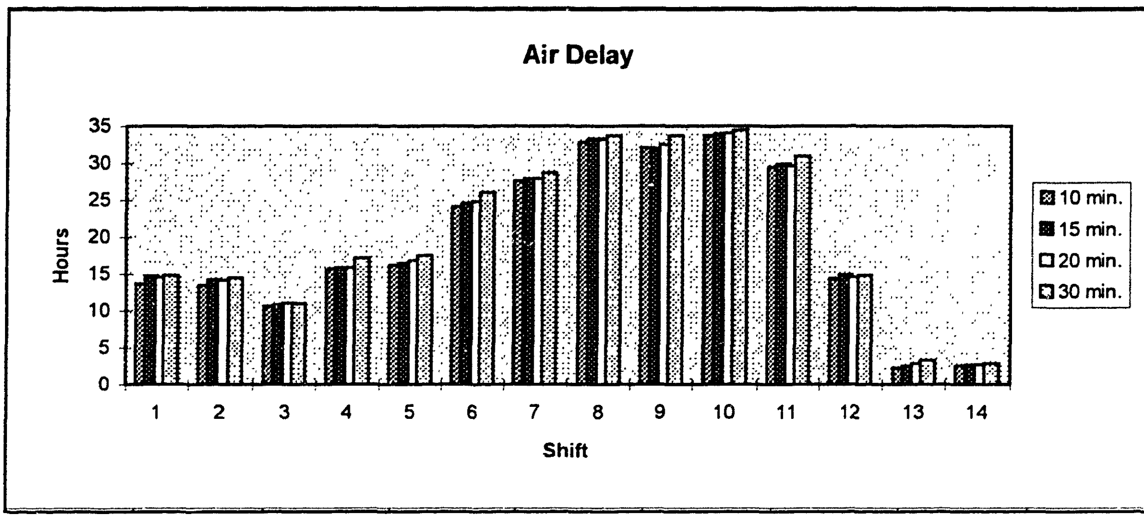

Fig. 4-14b: Case 1: Ground Delay for Various Values of T

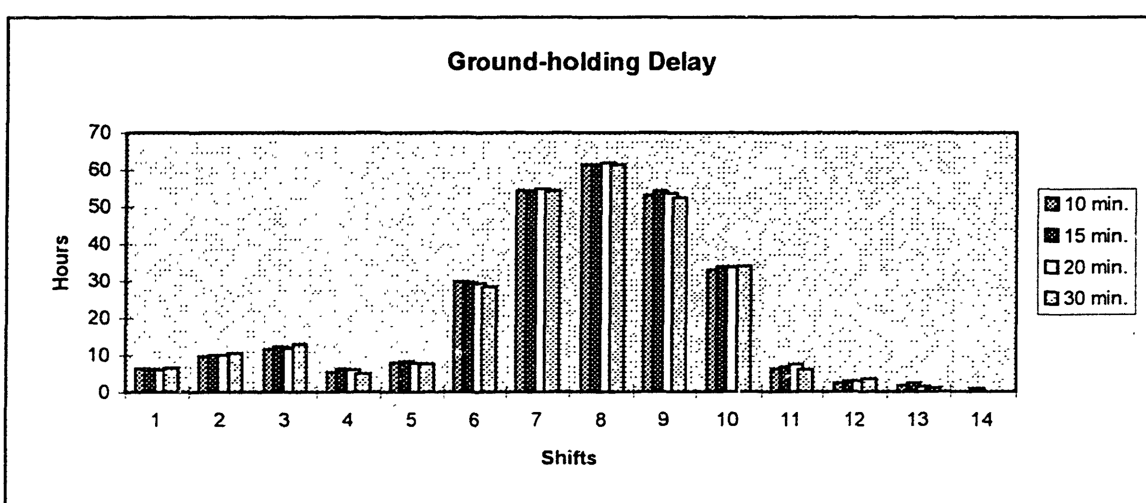

Fig. 4-15a: Case 2: Air Delay for Various Values of T

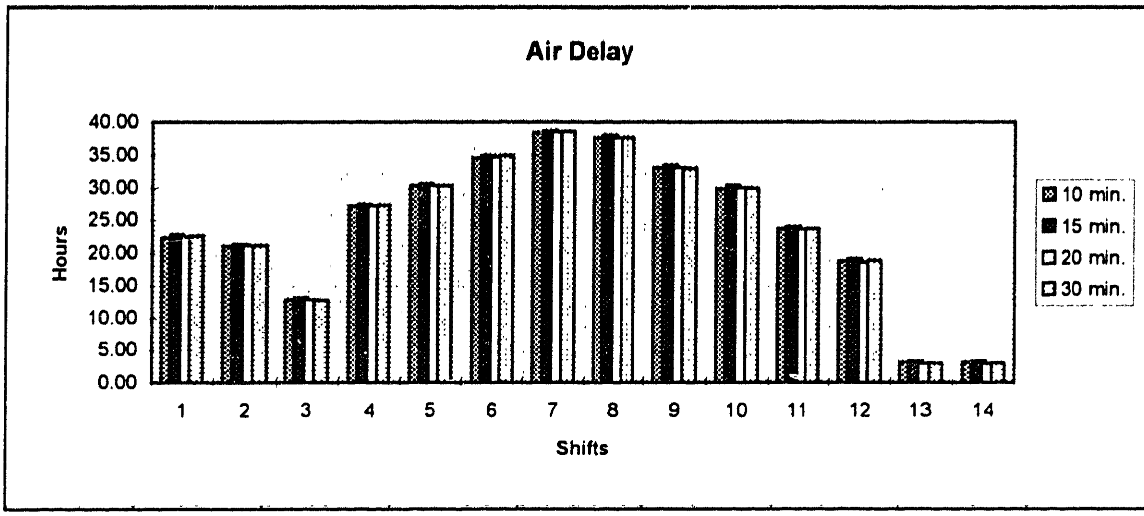

Fig. 4-15b: Case 2: Ground Delay for Various Values of T

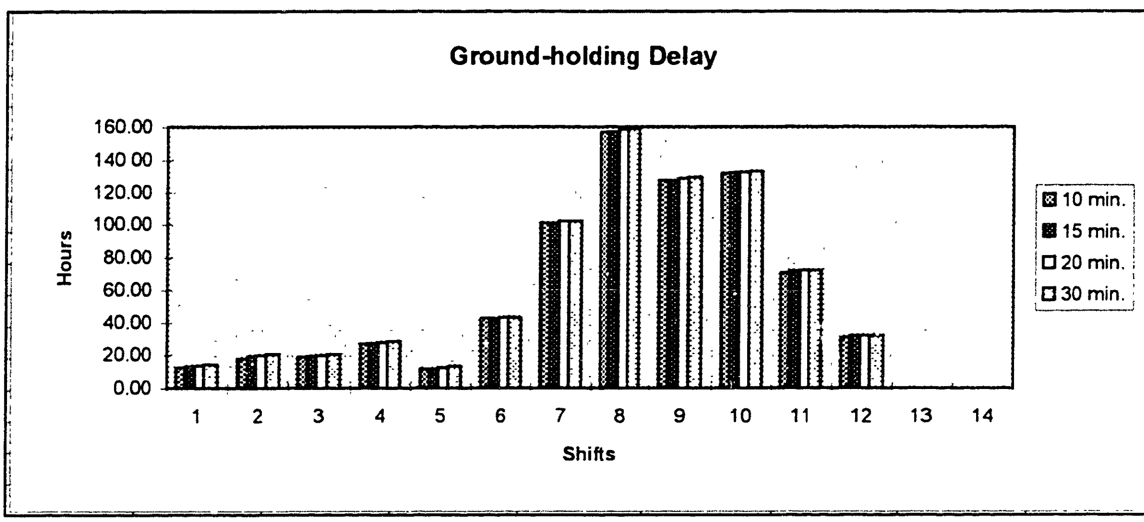

The above graphs show that the algorithm, in terms of the amount of air and ground delay, is not that sensitive to different values of T. Therefore, it is important to see the computation time associated with each T. Table 4-11 shows in the first row the matrix size of the problem associated with each value of T. The next three rows show, in seconds, the time to generate the matrix, the time to solve the problem, and the total time (sum of generation and solution time) for each value of T.

Table 4-11: Computation Time for Various Values of T

	30 min., T=34	20 min., T=51	15 min., T=68	10 min., T=102
Matrix Size	[142] x [836]	[210] x [1686]	[278] x [2825]	[414] x [5970]
Generation Time	.01	.3	.7	2.1
Solution Time	.6	1.5	3.0	9.3
Total time	.61	1.8	3.7	11.4

The above table shows that for T equal to 54 (30 minute periods), the problem is solved in less than a second. For T equal to 102 (10 minute periods) the computation time is a little over 11 seconds. These times were obtained for a Q equal to 3 and a time horizon of 17 hours. As was observed in the previous section, a larger Q will mean an increase in the computation time, but a larger Q may be more realistic. The time horizon of 17 hours seems reasonable, though. Thus, taking into account that Q may be larger than 3, dividing the time horizon into 15 minute periods appears to be an approach that will give good solutions quickly (as we saw in the previous section). However, because the algorithm does not appear to be that sensitive to T and can still be solved in a relatively short time, it seems that the decision on how to discretize the time horizon can be a flexible one. If an airport has a flat demand profile, then larger time period lengths may be used because the demand is more uniform. An airport with a demand profile that is more volatile and has many peaks and dips may wish to use smaller time period lengths.

4.4 Sensitivity to Cost Ratio

The purpose of this section is to show that the cost ratio of air delay per hour to ground delay per hour can influence what kind of ground-holding policy will be used. A high ratio of air delay cost to ground delay cost will result in an emphasis on ground-holding. This is because the cost of air delay is much higher than that of ground delay and hence, most delays will be ground ones. When air and ground delay costs are close in value, then more delays will be absorbed in the air because there is no longer any great savings associated with ground delays.

In this section, we use the same example used in the second section of this chapter. The demand and capacity profiles from Logan are used. Q equals 3 and T equals 68. For case 1, uncertainty in time, a two-hour storm could arrive at period 17, 21, or 25 and capacity will be reduced from 15 arrivals/15 min. to 7 arrivals/15 min. for these 2 hours. For case 2, uncertainty in capacity, a 4 hour storm is arriving at period 29 and during this time, capacity could be 15, 11, or 7 arrivals/15 min. We show three examples of each case, using different probability distributions. The three distributions used are: .333, .333, .333; .2, .3, .5; and .1, .1, .8, corresponding to the high, medium, and low capacity, respectively. For each distribution and each case, we find delays and costs for five cost ratios. These ratios are of the cost of air delay/hr. to the cost of ground delay/hr. and are as follows: \$1000/hr : \$200/hr. (5:1); \$1000/hr. : \$400/hr. (5:2); \$1000/hr. : \$600/hr. (5:3); \$1000/hr. : \$800/hr. (5:4); and \$1000/hr. : \$1000/hr. (5:5). The results from these runs can be seen in Tables 4-12a, 4-12b, and 4-12c for case 1 and in Tables 4-13a, 4-13b, and 4-13c for case 2. The top row indicates the ratio and in each column, the air delay, ground delay, total delay and the associated costs of each are shown. Delays are in hours and costs are in dollars.

Table 4-12a: Case 1: Cost Ratio for Probability Distribution .333, .333, .333

	5:1	5:2	5:3	5:4	5:5
Air Delay	7.99 hrs.	21.4 hrs.	21.4 hrs.	29.39 hrs.	30.89 hrs.
Ground Delay	53.75 hrs.	13.5 hrs.	13.5 hrs.	1.5 hrs.	0 hrs.
Total Delay	61.74 hrs.	34.9 hrs.	34.9 hrs.	30.89 hrs.	30.89 hrs
Air Cost	$7,992	$21,395.25	$21,395.25	$29,387.25	$30,885.75
Ground Cost	$10,750	$5,400	$8,100	$1,200	$0
Total Cost	$18,742	$26,795.25	$29,495.25	$30,587.25	$30,885.75

Table 4-12b: Case 1: Cost Ratio for Probability Distribution .2, .3, .5

	5:1	5:2	5:3	5:4	5:5
Air Delay	5.55 hrs.	6.1 hrs.	25.73 hrs.	25.73 hrs.	33.93 hrs.
Ground Delay	53.75 hrs.	51 hrs.	11.75 hrs.	11.75 hrs.	1.5 hrs.
Total Delay	59.30 hrs	57.1 hrs.	37.48 hrs.	37.48 hrs.	35.43 hrs.
Air Cost	$5,550	$6,100	$25,725	$25,725	$33,925
Ground Cost	$10,750	$20,400	$7,050	$9,400	$1,500
Total Cost	$16,300	$26,500	$32,775	$35,125	$35,425

Table 4-12c: Case 1: Cost Ratio for Probability Distribution .1, .1, .8

	5:1	5:2	5:3	5:4	5:5
Air Delay	2.68 hrs.	3.03 hrs.	3.03 hrs.	6.43 hrs.	42.25 hrs.
Ground Delay	51 hrs.	49.25 hrs.	49.25 hrs.	45 hrs.	1.5 hrs.
Total Delay	53.68 hrs.	52.28 hrs.	52.28 hrs.	51.43 hrs.	43.75 hrs.
Air Cost	$2,675	$3,025	$3,025	$6,425	$42,250
Ground Cost	$10,200	$19,700	$29,550	$36,000	$1,500
Total Cost	$12,875	$22,725	$32,575	$42,425	$43,750

Table 4-13a: Case 2: Cost Ratio for Probability Distribution .333, .333, .333

	5:1	5:2	5:3	5:4	5:5
Air Delay	0 hrs.	91.82 hrs.	91.82 hrs.	117.13 hrs.	117.13 hrs.
Ground Delay	313.75 hrs.	38 hrs.	38 hrs.	0 hrs.	0 hrs.
Total Delay	313.75 hrs.	129.82 hrs.	129.82 hrs.	117.13 hrs.	117.13 hrs.
Air Cost	$0	$91,824.75	$91,824.75	$117,132.75	$117,132.75
Ground Cost	$62,750	$15,200	$22,800	$0	$0
Total Cost	$62,750	$107,024.75	$114,624.75	$117,132.75	$117,132.75

Table 4-13b: Case 2: Cost Ratio for Probability Distribution .2, .3, .5

	5:1	5:2	5:3	5:4	5:5
Air Delay	0 hrs.	0 hrs.	137.88 hrs.	142.28 hrs.	168.28 hrs.
Ground Delay	313.75 hrs.	313.75 hrs.	38 hrs.	32.5 hrs.	0 hrs.
Total Delay	313.75 hrs.	313.75 hrs.	175.88 hrs.	174.78 hrs.	168.28 hrs.
Air Cost	$0	$0	$137,875	$142,275	$168,275
Ground Cost	$62,750	$125,500	$22,800	$26,000	$0
Total Cost	$62,750	$125,500	$160,675	$168,275	$168,275

Table 4-13c: Case 2: Cost Ratio for Probability Distribution .1, .1, .8

	5:1	5:2	5:3	5:4	5:5
Air Delay	0 hrs.	0 hrs.	0 hrs.	4.2 hrs.	254.8 hrs.
Ground Delay	313.75 hrs.	313.75 hrs.	313.75 hrs.	308.5 hrs.	0 hrs.
Total Delay	313.75 hrs.	313.75 hrs.	313.75 hrs.	312.7 hrs.	254.8 hrs.
Air Cost	$0	$0	$0	$4,200	$254,800
Ground Cost	$62,750	$125,500	$188,250	$246,800	$0
Total Cost	$62,750	$125,500	$188,250	$251,000	$254,800

For both case 1 and case 2, the above tables show that when the cost of ground delay is much cheaper than that of air delay, most of the delay will then be ground. But as the ground delay cost approaches that of air delay, most of the delay becomes air delay because ground-holding no longer represents high savings in cost. A solution with almost all of the delay as air delay is very similar to the passive strategy, which assigns all delay as air delay. The passive strategy gives the minimum amount of delay. That is why for each distribution, the total delay decreases as the ratio goes from 5:1 to 5:5.

In both cases, the third capacity scenario represents the worst case. For case 1, the third scenario represents a drop in capacity reduction during a higher demand period than the first and second scenarios. For case 2, the third scenario has the lowest capacity during the period of reduction. This means that as the probability of the third scenario occurring increases, more weight is given to this scenario and if it occurs, there will be more delay. As can be seen from the tables, the total delay is greater for probability distribution .1,.1, .8 then for distribution .333, .333, .333. The only exception to this is for case 1 and ratio 5:1. For this combination, the opposite occurs. This is because almost all delay is ground delay for ratio 5:1 and when there is more uncertainty about which particular scenario will occur, then there is a higher chance that available landing capacity is not being utilized. But as the probability of the third scenario increases, this scenario becomes the one that is most likely to occur and it is less likely that landing capacity is being wasted. This is also why for probability distribution .1, .1, .8, there is not as much discrepancy in total delay between ratio 5:1 and 5:5 as there is for distribution .333, .333, .333. This last observation is true for both cases.

The results of this section show that the ratio of air to ground cost chosen can influence the solution. The impact of this ratio on the solution can vary for different probability distributions and different problems, but it is still an important factor in determining the amount of ground-holding done. A high ratio of air to ground cost will result in a great deal of ground-holding, whereas a low ratio of air to ground cost means most of the delay will be air delay. It is important to keep this in mind when determining a cost ratio to use in the ground-holding problem. The ratio chosen can influence what kind of ground-holding policy will be used.

4.5 Comparison to Other Strategies

In this section, we compare the simplified algorithm developed in this thesis (for simplicity, we will call it the GHP algorithm) to three other strategies of air traffic flow management for landings at an airport. The Passive strategy is one where flights depart from the airports of origin on schedule and all delays are taken in the air. The Most Likely strategy chooses the probabilistic capacity scenario with the highest probability as the only capacity scenario and available capacity is then assigned on a FCFS basis, with all delays being ground-holding delays. The Expected strategy comes up with one capacity scenario by taking the expected value of the probabilistic capacity scenarios.

For the case of uncertainty in time, we use the same example as the second and fourth section of this chapter. We are looking at Logan over a 17 hr. day that starts at 7:01 AM and ends at midnight. We break the time horizon into 15 minute periods, to obtain a T equal to 68. Capacity will be reduced from 60 arrivals/hr. (15 arrivals/15 min.) to 30 arrivals/hr. (7 arrivals/15 min.) for a length of two hours. This reduction could begin at time period 17, 21, or 25. Each of these possible start times represents a capacity scenario. We run this example twice, using a different probability distribution in each case. The first run has a uniform distribution and the probabilities associated with each scenario is 1/3. The second run assigns the probabilities .8, .1, and .1 to the capacity scenarios with capacity reduction start time at periods 17, 21, and 25.

For the case of uncertainty in capacity, we also use the same example as described in the second and fourth section of this chapter. We know that between periods 29-44, capacity could be 15 arrivals/15 min., 11 arrivals/15 min., or 7 arrivals/15 min. We do two runs of this example, each with a different probability distribution. The two distributions are (associated with 15, 11, and 7 arrivals/15 min., respectively): 1/3, 1/3, 1/3 and .1, .1, .8.

Tables 4-14a and 4-14b show the delays and associated costs under the different scheduling strategies for the two probability distributions for case 1. Case 2 results are shown in Tables 4-15a and 4-15b. For the Most Likely strategy, under the probability distribution of 1/3, 1/3, 1/3, each scenario has equal chance of being the most likely. We apply the Most Likely strategy first to scenario 1, then to scenario 2, and lastly, to scenario 3.

Table 4-14a: Case 1: Probability Distribution: 1/3, 1/3, 1/3

	GHP Alg.	Passive	Most Likely (scenario 1)	Most Likely (scenario 2)	Most Likely (scenario 3)	Expected
Air Delay	21.4 hrs.	30.89 hrs.	53.5 hrs.	24.67 hrs.	10.08 hrs.	30.75 hrs.
Ground	13.5 hrs.	0 hrs.	22.5 hrs.	21 hrs.	49.25 hrs.	25 hr.
Total Delay	34.9 hrs.	30.89 hrs.	76 hrs.	45.67 hrs.	59.33 hrs.	31 hrs.
Air Cost	$21,395.25	$30,885.75	$53,500	$24,670	$10,080	$30,750
Ground Cost	$8,100	$0	$13,500	$12,600	$29,550	$150
Total Cost	$29,495.25	$30,885.75	$67,000	$32,270	$39,630	$30,900

Table 4-14b: Case 1: Probability Distribution: .8, .1, .1

	GHP Alg.	Passive	Most Likely (profile 1)	Expected
Air Delay	18.2 hrs.	25.03 hrs.	16.05 hrs.	17.83 hrs.
Ground Delay	8.5 hrs.	0 hrs.	22.5 hrs.	35.25 hrs.
Total Delay	26.7 hrs.	25.03 hrs.	76 hrs.	53.08 hrs.
Air Cost	$18,200	$25,025	$16,050	$17,830
Ground Cost	$5,100	$0	$13,500	$21,150
Total Cost	$23,300	$25,025	$29,550	$38,980

Table 4-15a: Case 2: Probability Distribution: 1/3, 1/3, 1/3

	GHP Alg.	Passive	Most Likely (scenario 1)	Most Likely (scenario 2)	Most Likely (scenario 3)	Expected
Air Delay	91.82 hrs.	117.13 hrs.	117.25 hrs.	91.92 hrs.	0 hrs.	91.92 hrs.
Ground	38 hrs.	0 hrs.	0 hrs.	38 hrs.	313.75 hrs.	38 hrs.
Total Delay	129.82 hrs.	117.13 hrs.	117.25 hrs.	129.92 hrs.	313.75 hrs.	129.92 hrs.
Air Cost	$91,824	$117,133	$117,250	$91,920	$0	$91,920
Ground Cost	$22,800	$0	$0	$22,800	$188,250	$22,800
Total Cost	$114,624	$117,133	$117,250	$114,720	$188,250	$114,720

Table 4-15b: Case 2:Probability Distribution: .1, .1, .8

	GHP Alg.	Passive	Most Likely (scenario 3)	Expected
Air Delay	0 hrs.	254.8 hrs.	0 hrs.	266.4 hrs.
Ground Delay	313.75 hrs.	0 hrs.	313.75 hrs.	230.5 hrs.
Total Delay	313.75 hrs.	254.8 hrs.	313.75 hrs.	496.9 hrs.
Air Cost	$0	$254,800	$0	$266,400
Ground Cost	$188,250	$0	$188,250	$138,300
Total Cost	$188,250	$254,800	$188,250	$404,700

The above tables show that for these particular examples, the ground-holding algorithm developed in this thesis always gives the lowest total cost and a total delay that is close to the minimum possible. The other strategies are much more erratic. Sometimes, they give solutions with costs that are close to the lowest and other times, their solutions have the highest costs. The performance of these scheduling strategies depend on the demand and capacity profiles used and the associated probability distribution. The GHP algorithm is more consistent and performs better under different situations.

Chapter 5

Conclusions and Future Research

5.1 Conclusions

In this section, we summarize the conclusions drawn in this thesis.

The cost of air delay is, in practice, always higher than the cost of ground delay. Whenever demand exceeds capacity, a ground-holding policy may thus be beneficial, in terms of cost savings. Times of high demand in relation to capacity will use more ground-holding and will thus, have higher savings. An airport that does not have much variability in its capacity nor a high demand may not need a ground-holding algorithm. Also, some airports do not have persistent high demand all day. Ground-holding may then only be needed during these peak travel times.

Our algorithm is sensitive to Q, the number of probabilistic capacity scenarios used. It is important to model the true probability distribution as closely as possible in order to obtain good results. This may mean that Q could be as large as 10. When we try to use smaller than necessary values of Q, problems arise in determining how many and which scenarios to use. Different values of Q and different combinations of scenarios give different results. We found, however, that if the number of time periods is around 68, then a Q of 10 can be used and the problem can be solved in a relatively short time. With T, the number of time periods, equal to 68 a 17 hour day, divided into 15 minute periods, can be represented. This is an acceptable assumption because demand at night is low and flights at night usually have no delays in landing.

The algorithm does not appear to be particularly sensitive to T. Varying T did not appear to affect the results much. This gives flexibility to the user of the algorithm in determining what value of T to use. Depending on an airport's demand and capacity profiles, different T values may be chosen. An airport with high peaks of demand and more variability may require smaller time period lengths (a larger T), whereas an airport with a uniform or flat demand profile and a stable capacity could use larger time period lengths (a smaller T).

The ratio of the cost of air delay per time period to the cost of ground delay per time period can influence what kind of ground-holding policy will be implemented. A high ratio of air to ground cost will call for more ground-holding than a low ratio. Therefore, this ratio should be determined carefully.

A comparison of the ground-holding algorithm with other flow control strategies indicated that the algorithm offered the lowest total cost and close to the minimum amount of delay possible. The ground-holding algorithm was also more consistent in its performance under various situations than the other strategies.

The results of this thesis seem to indicate that a carefully developed ground-holding policy will be beneficial in reducing delay costs, especially in situations where the demand exceeds the capacity by a great deal. When using the ground-holding algorithm, we need to be careful in choosing the number of probabilistic capacity scenarios used and the cost ratio chosen. This algorithm can solve large problems quickly and should be able to be used in practice. The GUI described in Chapter 3 shows that this algorithm, when combined with a computer program, becomes a decision support system that solves problems in real time, gives simple answers, and is user friendly.

5.2 Future Research

In this thesis, we have attempted to examine the potential use of a probabilistic optimization model to solve the ground-holding problem for a single airport. We have tried to find the parameters that had the greatest effect on the results and running time. An important outcome of the experiments was a better understanding of the sensitivity of Richetta's algorithm to the probability distribution of the airport capacity and to the number of capacity scenarios used.

We have found that it is better to use the true probability distribution, rather than an approximation, and that as far as running time is concerned, this is a practical approach. It would be beneficial, however, to do more research in this area. One possibility is to connect Richetta's algorithm to a simulation tool in order to generate capacity scenarios that model realistic situations. Now that we have seen that the algorithm's running time is small, we could also examine Richetta's algorithm, using different classes of aircraft, and run the same kind of experiments as in this thesis.

This thesis looked at the static model, which is a fair representation of an airport that makes one decision at the start of the day for a ground-holding policy. Many airports, though, are better represented by a dynamic model. Because our model can be solved quickly and gives good results, it can be solved many times during the day, thereby resembling a dynamic model. A logical next step for future research may be to perform a probabilistic analysis of a dynamic ground-holding model and then do a comparison between the dynamic model and the static model that is solved repeatedly through the day.

Air traffic flow is not the only area of transportation that the ground-holding strategy can be applied to. Any type of situation in which there are flows into one node that becomes capacitated could benefit from some kind of ground-holding policy.

Bibliography

[1] G. Andreatta and L. Brunetta (1995), *Multi-Airport Ground Holding Problem: A Computational Evaluation of Exact Algorithms*, Technical Report, Department of Pure and Applied Mathematics, University of Padova, Padova, Italy.

[2] G. Andreatta, L. Brunetta and G. Guastalla (1994), *Multi-Airport Ground Holding Problem: A Heuristic Approach Based on Priority Rules*, Technical Report, Department of Pure and Applied Mathematics, University of Padova, Padova, Italy.

[3] D. Bertsimas and S. Stock (1994), *The air traffic flow management problem with enroute capacities*, Operations Research Center Working Paper #3726-94, MIT.

[4] E. A. Boyd, R. Burlingame and K. Lindsay (1994), *An Algorithm for the Management of National Air Traffic*, Report INEN/OR/WP/10/5-94, Texas A&M University, College Station, Texas, USA.

[5] R. Burlingame, E. A. Boyd and K. Lindsay (1995), "Solving Large Integer Programs Arising from Air Traffic Flow Problems," in *Intelligent Scheduling Systems*, Donald E. Brown and W. T. Scherer (eds.), Kluwer Academic Publishers, Boston, MA, pp. 215-226.

[6] L. Epstein, A. Futer and L. Medvedovsky (1992), *Multiple Airport Scheduling*, ATMS Interim Report, Volpe National Transportation Systems Center, Cambridge, MA, USA.

[7] M. P. Helme (1992), "Reducing Air Traffic Delay in a Space-Time Network", *Proceedings of the 1992 IEEE International Conference on Systems, Man and Cybernetics*, Chicago, pp. 236-242.

[8] M. P. Helme (1994), *The Potential of Network Flow Models for Managing Air Traffic*, Report MTR 94W0000058, The MITRE Corporation, McLean, VA, USA.

[9] K. S. Lindsay, E. A. Boyd and R. Burlingame (1993), "Traffic Flow Management Modeling with the Time Assignment Model", *Air Traffic Control Quarterly, 1*, no. 3, pp. 255-276.

[10] L. Navazio and G. Romanin-Jacur (1995), *Multiple Connections Multi-Airport Ground Holding Problem: Models and Algorithms*, Working Paper, Department of Electronics and Informatics, University of Padova, Padova, Italy.

[11] W. Philipp and F. Gainche (1994), "Air Traffic Flow Management in Europe", in *Advanced Technologies for Air Traffic Flow Management*, H. Winter and H. G. Nusser, editors, Springer-Verlag, Berlin, pp. 64-106.

[12] O. Richetta and A. R. Odoni (1993), "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control", *Transportation Science, 27*, pp. 228-238.

[13] O. Richetta and A. R. Odoni (1994), "Dynamic Solution to the Ground-Holding Policy Problem in Air Traffic Control", *Transportation Research A, 28*, pp. 169-185.

[14] O. Richetta (1991), *Ground Holding Strategies for Air Traffic Control Under Uncertainty*, Operations Research Center Technical Report No. 198, MIT, Cambridge, MA.

[15] M. Terrab (1990), *Ground Holding Strategies in Air Traffic Control*, Operations Research Center Technical Report No. 196, MIT, Cambridge, MA.

[16] M. Terrab and A. R. Odoni (1993), "Strategic Flow Control on an Air Traffic Network", *Operations Research, 41*, pp. 138-152.

[17] M. Terrab and S. Paulose (1993), *Dynamic, Strategic and Tactical Air Traffic Flow Control*, RPI Technical Report.

[18] P. B. Vranas (1992), *The Multi-Airport Ground Holding Problem in Air Traffic Control*, PhD Thesis, Operations Research Center, MIT, Cambridge, MA.

[19] P. B. Vranas, D. Bertsimas and A. R. Odoni (1994a), "The Multi-Airport Ground-Holding Problem in Air Traffic Control", *Operations Research, 42*, No. 2, pp. 249-261.

[20] P. B. Vranas, D. Bertsimas and A. R. Odoni (1994b), "Dynamic Ground-Holding Policies for a Network of Airports", *Transportation Science*, 28, No. 4, pp. 275-291.

THESIS PROCESSING SLIP

FIXED FIELD: ill ______ name ______

index ______ biblio ______

► COPIES Archives Aero Dewey Eng Hum

Lindgren Music Rotch Science

TITLE VARIES ► ☑ in degree book, first word is "Probabilistic

NAME VARIES: ► ☐

IMPRINT (COPYRIGHT)

► COLLATION: 79 l

► ADD. DEGREE: ______ ► DEPT.: ______

SUPERVISORS:

NOTES:

cat'r: date:

page: F62

► DEPT: O.R.

► YEAR: 1997 ► DEGREE: M.S.

► NAME: SHEEL, Minakshi