A Model to Design a Stochastic and Dynamic Ground Delay Program Subject to Non-Linear Cost Functions

by

Michael J. Hanowsky

B.S. Industrial Engineering and Operations Research – University of California, 2000
S.M. Operations Research – Massachusetts Institute of Technology, 2006
M.S. Transportation – Massachusetts Institute of Technology, 2006

Submitted to the Engineering Systems Division in Partial Fulfillment of the Requirements for the Degree of

DOCTOR OF PHILOSOPHY IN ENGINEERING SYSTEMS

at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
June 2008

Signature of Author ....
Engineering Systems Division
May 20, 2008

Certified By ....
Cynthia Barnhart, Professor of Civil and Environmental Engineering and Engineering Systems and Associate Dean for Academic Affairs, School of Engineering
Thesis Committee Chair

Certified By ....
Amedeo R. Odoni, Professor of Civil and Environmental Engineering and Aeronautical and Astronautical Engineering
Thesis Supervisor

Certified By
Joseph M. Sussman, JR East Professor of Civil and Environmental Engineering and Engineering Systems
Thesis Committee Member

Accepted By..
Richard C. Larson, Mitsui Professor of Engineering Systems and Civil and Environmental Engineering
Chair, ESD Education Committee

A Model to Design a Stochastic and Dynamic Ground Delay Program
Subject to Non-Linear Cost Functions
by

Michael J. Hanowsky

Submitted to the Engineering Systems Division
On May 20th, 2008 in partial fulfillment of the
Requirements for the Degree of Doctor of Philosophy in
Engineering Systems

ABSTRACT

When inclement weather reduces the arrival capacity of a busy metropolitan airport, it may lead to significant airborne delays. Delaying aircraft in the air consumes additional fuel, increases overall air traffic congestion, and may lead to costly flight diversions. As a result, during periods of inclement weather, the FAA may implement a Ground Delay Program (GDP) to proactively delay flights on the ground before they depart and reduce the possibility of future airborne delays. However, in order to assign ground delays to flights, a GDP must be implemented before they depart, at a time when the future airport arrival capacity may be uncertain.

This dissertation discusses two analyses in regards to the design of a GDP. The first analysis proposes a model that solves for the optimal assignment of ground delay to aircraft for a stochastic and dynamic forecast of the airport arrival capacity, with non-linear delay cost functions, and a capacity of the airborne arrival queue. This model is applied to several hypothetical examples and, in comparison to prior models from the literature, identifies solutions with a lower total expected cost, a smaller maximum observed arrival queue, or both.

The second analysis compares the salience, or importance, of various stakeholder groups to their roles in the design of a GDP in practice. Passengers, in particular, are shown to be an important, but under-represented stakeholder group. A second model is proposed that solves for an assignment of ground delay that minimizes the total passenger delay cost. A comparison of these results to those of the first model show that the total cost of delays to passengers could be reduced by more than 30% if the FAA were to directly consider the cost of delays to passengers during the design of a GDP.

Thesis Supervisor: Amedeo Odoni
Title: Professor of Aeronautics and Astronautics
Professor of Civil and Environmental Engineering

Acknowledgements

I am grateful for the thoughts, advice, and guidance provided by my committee and, especially, Rick Oiesen and Dr. Eugene Gilbo at the John A. Volpe National Transportation Systems Center.

In addition, I would also like to thank Bill Leber, Ken Howard, Roger Beatty, Tim Matuszewski, and the employees of Northwest Airlines and United Airlines, who provided me with perspective and insight.

And, most of all, I am indebted to the support of my friends and family.

Table of Contents

Chapter 1 Introduction

The air transportation industry in the United States has grown significantly in the past 20 years. Deregulation of commercial air services, changes in the business environment, and new materials and manufacturing technologies have led to a greater number and variety of aircraft in the skies than ever before, from the proliferation of small corporate air taxi services to the use of the much-publicized Airbus A-380. As the industry has grown, so has the use of the nation's airspace, such that the scheduled demand for the use of various air routes and runways may meet or exceed the planned capacity and, therefore, result in significant delays to aircraft, crews, and passengers.

One of the responsibilities of the Federal Aviation Administration (FAA[1]) is the management of the flow of air traffic through the National Airspace System (NAS). In cases where excess demand is a chronic condition, such as at busy metropolitan airports in New York and Chicago, the FAA may reduce demand by placing limits on the number of scheduled flight operations, or increase capacity by redesigning the airspace or working with airports to build additional runways. However, for situations in which the excessive demand is unplanned, such as when inclement weather reduces the capacity of a sector or airport, it may not be possible to increase the capacity of the system. In these cases, the FAA may implement various air traffic flow management (ATFM) techniques to reroute or delay flights in the air or on the ground.

[1] A list of acronyms is provided in the appendix

Inclement weather is of particular concern because it can result in substantial delay costs. Not only may weather-related reductions in NAS capacity be significant, but the timing and severity of the decrease may also be subject to uncertainty. Busy metropolitan airports are especially prone to weather-related delays and congestion, where poor weather, such as changes in wind speed and direction, visibility, and precipitation, can reduce the airport arrival capacity by one-half, or more. Furthermore, as there may be limited buffer capacity during nominal operating conditions, a temporary reduction in the arrival capacity of one airport may lead to delays throughout the system that persist for hours.

To mitigate the congestion and other costs of delays due to insufficient arrival capacity at an airport, the FAA may implement a Ground Delay Program, or GDP. The objective of a GDP is to avoid extensive airborne delay by proactively delaying flights on the ground before they depart. Ground delay is preferable to airborne delay because delays on the ground may avoid the excess fuel consumption, flight diversions, and overall air traffic congestion that are associated with airborne delays. However, in order to assign ground delays, a GDP must be initiated before flights depart, at which time the future capacity of the destination airport may be uncertain.

This uncertainty presents a key tradeoff for the design of a GDP. The earlier that a GDP is initiated, the more flights that are eligible to receive ground delay, but the greater the uncertainty in the arrival capacity of the airport (Figure 1-1). If a program is initiated before the airport arrival capacity is known, then there is a risk that it may assign more ground delay than is necessary to offset the actual reduction in arrival capacity. This additional, or unrecoverable, ground delay not only has a direct cost, but may also result in a cascade of delays throughout the day, as the aircraft to which it assigns ground delay may have little buffer in their daily flight schedules.

On the other hand, if too little ground delay is assigned, or if the implementation of a GDP is postponed until more information is available, extensive airborne delays may result. Airborne delays may be very costly because, as aircraft run low on fuel, they may divert to an alternate airports. International and long distance flights are particularly susceptible because they may not carry sufficient fuel to endure an airborne delay of more than a few minutes. Furthermore, airborne delays can result in air traffic congestion, with may result in delays for flights that are flying to other, unconstrained airports.

Figure 1-1 The tradeoff between control and information over time

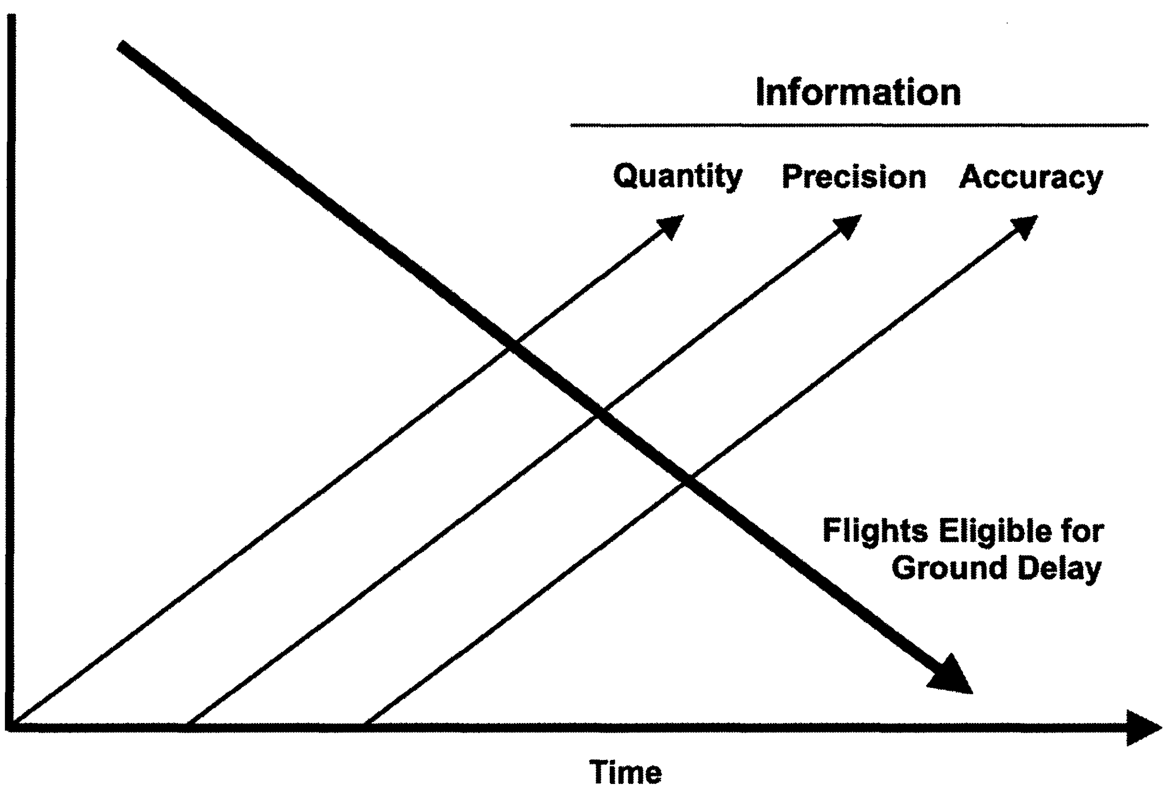

This dissertation is motivated by the question of how to assign ground delays in order to minimize the costs due to a weather-related reduction in the arrival capacity of an airport. The question may be characterized by three key qualities of the problem. First, decisions must often be made as based upon *uncertain* forecasts of the airport arrival capacity and demand. Second, the decision is *dynamic*; not only does the uncertainty of the forecast decrease over time, but the option to assign delay to flights also decreases as aircraft depart (Figure 1-1). Furthermore, decisions made at a given time may influence the opportunity to make additional decisions at later times. And third, a GDP will have wide-ranging impacts on a variety of stakeholders and the costs of delays to these stakeholders may not be uniform. Thus, two important questions regarding the design of a GDP are:

1. Who or what should be considered during the design of a GDP and
2. What is the optimal assignment of ground delay?

This chapter provides a brief overview of the design of GDPs in practice and in the academic literature and then summarizes the contribution of this thesis.

Section 1.1 The Design of Ground Delay Programs In Practice

In practice, a GDP is designed by an air traffic manager at the FAA Air Traffic Control System Command Center (ATCSCC) in Herndon, VA. The traffic manager is responsible for three key decisions: 1) whether or not to implement a GDP, 2) which flights are eligible to receive ground delay, and 3) the choice of planned airport arrival rates, which reflects the opinion of traffic manager as to what the future arrival capacity of the airport will be. In making these decisions, the traffic manager will consider one or more forecasts of the airport arrival capacity and a list of flights that are scheduled to arrive at the airport. In addition, the traffic manager may consult with representatives from the commercial airlines, as well as other traffic managers from local air traffic control facilities, such as the airport terminal radar approach control facility (TRACON) or the air route traffic control centers (ARTCCs).

Once the decision to implement a GDP has been made, the traffic manager assigns ground delays using the Ration-By-Schedule (RBS) algorithm. RBS assigns ground delays to eligible flights in a manner consistent with a first-scheduled-first-served order (FSFS). The outcome of the RBS algorithm is dependant on the choices made by the traffic manager; only those flight that are included in the GDP may be assigned ground delay and the amount of ground delay that is assigned depends on the chosen arrival capacity rate.

A key quality of the assignment of delays is that from the perspective of the traffic manager and the RBS algorithm, all flights are considered to have an equal priority. Other factors, such as the number of passengers, and passenger, crew, and equipment connections are not explicitly considered. To account for additional considerations, the ground delays that are initially assigned by the traffic manager (using RBS) may be revised by the airlines using various mechanisms, such as flight swaps and slot-credit-substitution. A key to each of these mechanisms is that they allow the ground delays to be revised to reflect information that is available to the airlines but will not increase the arrival demand as controlled by the GDP for any period of time. Additional discussion on the design of GDPs in practice may be found in Wambsganss (2001) and Chang et al. (2001).

The assignment of ground delay in practice is based on a qualitative evaluation of the problem. First, the decisions to assign of ground delay to aircraft are based on the experience and expertise of the traffic manager. As a result, given the same set of inputs, two different traffic managers might design the GDP differently. Second, the forecast uncertainty of the airport arrival capacity is not quantified by the design of a GDP; one arrival acceptance rate is chosen for the purposes of planning the GDP. In actuality, however, there may be one or more different forecasts of the future arrival capacity, each of which has some likelihood of occurring. The choice that is used for the design of the GDP may be the most likely capacity, but that does not guarantee that it will occur. Third, the design process does not quantify the various costs of choosing an incorrect arrival capacity rate. As a result, the design of a GDP in practice reflects a subjective evaluation of the available information and might be sub-optimal from the perspective of the system.

Section 1.2 The Design of Ground Delay Programs in the Academic Literature

In the academic literature, the design of a GDP is represented as the single airport ground holding problem (SAGHP). Various models have been proposed in the SAGHP literature that use mathematical optimization techniques to solve for an assignment of ground delay. One advantage of these models is that, subject to a set of input parameters and an objective function, the assignment of ground delay that is defined by the solution is both objective and optimal. Furthermore, the academic literature has also proposed the use of scenario trees to model stochastic and dynamic forecasts of the airport arrival capacity. Using these scenario trees, a series of models have been proposed to solve the SAGHP subject to a stochastic and dynamic arrival capacity forecast.

These stochastic/dynamic SAGHP models are advantageous because they are able to consider many of the characteristics of the problem of designing a GDP. However, as with any model, the value or utility of the solution depends on the degree to which the model and its input parameters represent the problem as it occurs in practice. For example, existing SAGHP models are limited to objective functions of a particular form and do not consider either the costs of flight diversions, which occur due to lengthy airborne flight delays, or limits that might be placed on the size of the airborne arrival queue by the traffic manager. As a result, the solutions to these models may permit more airborne delays than would be desired in practice by the traffic manager and might also result in more diversions, which are undesirable to both airlines and their passengers..

Section 1.3 Contributions of the Thesis

This thesis is presented in two parts. The first part of the thesis builds upon the wealth of existing models in the literature and proposes a new model to solve the stochastic/dynamic SAGHP. The second part explores the design of GDPs in practice by comparing the salience of various stakeholders to their roles in the design process.

Section 1.3.1 A New SAGHP Model

The first contribution of this thesis is a new optimization model that solves the SAGHP subject to a stochastic and dynamic forecast of the airport arrival capacity. This new model, the ECM, improves upon those in the literature by also considering a wider range of delay cost functions and a capacity of the airborne arrival queue. The objective function of the ECM assumes that the cost of the delays may be a non-linear function of the duration of airborne, ground, and cumulative delay to each flight, where cumulative represents the sum of airborne and ground delays. It is assumed that non-linear functions may more accurately represent costs such as flight diversions and missed connections than functions that are strictly linear. The ECM allows the traffic manager to specify a maximum limit, or capacity, of the airborne arrival queue permitted by the optimal solution.

Two sets of experiments are conducted using the model. The first set examines the sensitivity of the solution to the model to changes in various input parameters, with a focus on those parameters that are unique to the ECM as compared to other stochastic and dynamic SAGHP models. The analysis indicates that it is possible to identify certain "critical" decision times at which the decisions regarding the implementation of a GDP and the assignment of ground delay should be made.

The sensitivity analysis also identifies a relationship between these objectives and the times at which ground delays are assigned. For the given example, the total expected cost of the solution to the ECM is sensitive to the time at which ground delays are revised subject to perfect information. However, the cost is much less sensitive to the time at which ground delays are assigned subject to an initial, uncertain forecast. On the other hand, the ability to identify solutions with smaller arrival queues is much more sensitive of the time of the initial GDP than to subsequent revisions.

The second set of experiments is an arrival study case study analysis, which solves the SAGHP for a variety of arrival capacity forecasts and with previous models from the SAGHP literature. The analysis shows that not only can the ECM be applied to such forecasts, but that it also results in solutions that are preferable to those of the previous models. For each example, the solution to the ECM either results in an equal or lower cost, a smaller maximum observed arrival queue length, or both.

Section 1.3.2 A Stakeholder Approach to the Design of a GDP

As an extension of the SAGHP, an analysis is presented which examines the salience of different stakeholder groups in regards to the design of a GDP in practice. Various stakeholder groups are described and classified according to their power, legitimacy, and urgency in regards to the design of a GDP. Airlines and the FAA are identified as being the most salient stakeholders, which is in accordance with their roles in the GDP design process. However, passengers, are also identified as an important, but under-represented stakeholder in the design process. Although this result may not seem surprising (to passengers), previous SAGHP models have not considered how delays might affect passengers.

This dissertation presents a final analysis to demonstrate how passengers might be considered during the design of a GDP. A second SAGHP model is proposed that assigns ground delays to minimize the total passenger delay cost. This formulation is applied to a simple example and the results show that the delay cost to passengers might be significantly reduced if the traffic manager were to explicitly consider the cost of delay to passengers during the initial design of a GDP.

Chapter 2 A Review of the Literature

A Ground Delay Program (GDP) is a tool used by the FAA to manage the arrival demand at a capacity-constrained airport, often as a response to inclement weather, which may reduce the arrival capacity of an airport significantly. GDPs intentionally delay aircraft on the ground prior to departure in order to avoid the possibility of future airborne delays. Ground delay is preferable to both the FAA and the airlines that use the NAS because it is less costly and it reduces airborne congestion. However, in order to delay flights on the ground, a GDP must be initiated before they depart, at which time the future arrival capacity of the airport may be uncertain.

In the academic literature, the decision to assign ground delay to aircraft is represented as the ground hold problem (GHP). The objective of the models that solve the GHP is to identify an assignment of ground delay to aircraft that minimizes the total cost of ground and airborne delay across all flights. Within the GHP, the models that are most closely related to the design of a GDP in practice are those that assign delay to flights that are scheduled to arrive at a single airport. The development of these single-airport GHP (SAGHP) models has focused on the assignment of delay in consideration of stochastic and dynamic forecasts of the airport arrival capacity.

This chapter discusses the stochastic/dynamic SAGHP models in the literature and identifies an opportunity to improve how they consider the objectives of a GDP as designed in practice. First, §2.1 and §2.2 outline the GHP and discuss the use of optimization techniques by models in the literature to solve for the assignment of ground delay. Then, §2.3 and §2.4 discuss the SAGHP and opportunities to improve upon SAGHP models in the literature.

Section 2.1 A Description of the Ground Holding Problem

In the academic literature, the decision to hold aircraft on the ground as a means of air traffic flow management (ATFM) is represented as the ground holding problem (GHP). Although the origins of the GHP in the research literature may be traced as far back as thirty years[2], the first systematic treatment of the problem is by Odoni (1987). Odoni describes the use of ground delays as part of a more general air traffic flow management problem, in which the movement of aircraft through the NAS is managed by the FAA and subject to the capacity of various elements of the system, such as air routes and airports.

As described by Odoni, the decision to assign ground delays embodies several key characteristics:

1. The demand for and capacity of system elements, such as airport arrivals, are *stochastic*, or subject to uncertainty
2. Information regarding the demand and capacity and the ability to assign ground delay are *dynamic* and will evolve over time
3. The use of ground delays may be measured by its distributive effects on individual aircraft
4. Ground delays will also have aggregated effects on groups of aircraft, airlines, and other stakeholders

[2] Odoni (1987)

The review of GHP literature presented in this chapter highlights the use of optimization models to assign ground subject to these considerations, and focuses on models designed for the case of a single airport. For a more comprehensive (and excellent) review of the GHP literature, please refer to Hoffman et al. (2007) or Ball et al. (2006).

Section 2.2 An Overview of Optimization Models for the GHP

Following this initial description, numerous models have been proposed in the literature that solve for an assignment of ground delay to aircraft in consideration of various aspects of the GHP. A common approach has been to use mathematical optimization techniques to solve for an assignment of flight departure times that minimizes the sum of delay costs to each flight. Key differences between the models are the assumptions that each makes in regards to the characteristics of the problem; to date, no tractable model has been proposed that captures all of the aspects identified by Odoni.

Models in the GHP literature are often classified by the number of capacity-constrained NAS resources that are considered. Terrab and Odoni (1993) propose an integer programming (IP) formulation that assigns ground delays to aircraft that are scheduled to arrive at a single airport with deterministic capacity. The formulation minimizes the sum of the delay costs attributed to each flight, which assumes that the function of delay cost is known *a priori*. Hoffman and Ball (2000) propose an IP model for a single airport that considers additional banking constraints, representing how delays affect the operations of a hub airline.

Other models have been proposed that consider multiple resources. For example, Vranas et al. (2004) propose an IP formulation for a network of multiple airports with deterministic capacity. Mukherjee (2004) proposes an IP formulation that considers the capacities of multiple arrival fixes at a single airport. Bertsimas and Stock-Patterson (1998) propose an IP that considers the capacity of both multiple airports and en route sectors. The advantage of the multiple resource models is that they may capture the downstream effects of delays at a single airport, as well as the interaction between capacity constraints for multiple airports and en route sectors. However, the multiple resource models require more information, resulting in more variables and constraints, which may exceed computational resources.

Within the larger body of literature, one area of focus has been to explore how solutions to the GHP distribute delays among aircraft and airlines. Odoni (1987) suggests that strategies for optimizing the flow of aircraft may result in a systematic bias in favor of long-distance flights. Hanowsky (2007) demonstrates that an optimal assignment of ground delay to flights scheduled to arrive at a single airport assigns more delay to short-distance flights and, as a result, to the airlines that operate these flights. Additionally, Lulli and Odoni (2007) show that if multiple resources are considered, then a reduction in capacity to one resource may result in ground delays for flights that are not scheduled to use the constrained resource.

Furthermore, several models have been developed in the literature to incorporate the distribution of ground delay among flights into the objectives of the GHP. Vossen (2002) proposes a model for mitigating the bias in assigned delay due to flight exemptions. Mukherjee (2004) proposes an optimization model that uses a super-linear cost function to identify solutions that are more equitable. Mukherjee also demonstrates that it is possible to adjust the balance of equity vs. efficiency in the solution by varying the weight assigned to such a function.

Section 2.3 The Single Airport Ground Hold Problem

A GDP is one type of ATFM initiative in which the assignment of ground delay is specific to flights that are scheduled to arrive at a single capacity-constrained airport. As ground delays are assigned before these flights depart, the design of a GDP may be subject to an uncertain forecast of the airport arrival capacity. In the academic literature, the design of a GDP is represented by a sub-problem within the GHP that is referred to as the single-airport ground hold problem (SAGHP). Published research on the SAGHP has contributed to two general areas: (1) the development of quantitative models to represent uncertain and dynamic arrival capacity forecasts and (2) the use of optimization methods to solve the SAGHP in consideration of stochastic and/or dynamic capacity forecasts.

Section 2.3.1 A Model of a Stochastic and Dynamic Arrival Capacity Forecast

A principal input to the design of a GDP in practice is a planned arrival acceptance rate, or PAAR, which is the maximum rate at which an airport is expected to be able to accommodate arriving aircraft. Although the PAAR may be set to a nominal rate to represent regular operating conditions, unscheduled reductions in the acceptance rate – resulting from exogenous factors, such as changes in wind speed and direction, visibility, and precipitation – will often trigger the initiation of a GDP in practice.

Because atmospheric conditions may change over the course of a day, the forecast arrival capacity of an airport can be represented as a *profile* of PAARs that vary over time[3].

Terrab and Odoni (1993) and Richetta and Odoni (1993) represent an arrival rate forecast in the presence of uncertainty as a mutually exclusive and collectively exhaustive set of profiles, each of which has an associated likelihood of representing the actual airport arrival capacity. Let an arrival capacity *scenario* consist of the set of possible profiles that correspond to a given forecast. The profile that actually occurs is said to be *realized.* An important assumption regarding any arrival capacity scenario is that exactly one of the profiles in the scenario will ultimately be realized.

Figure 2-1 An example of an arrival capacity scenario forecast

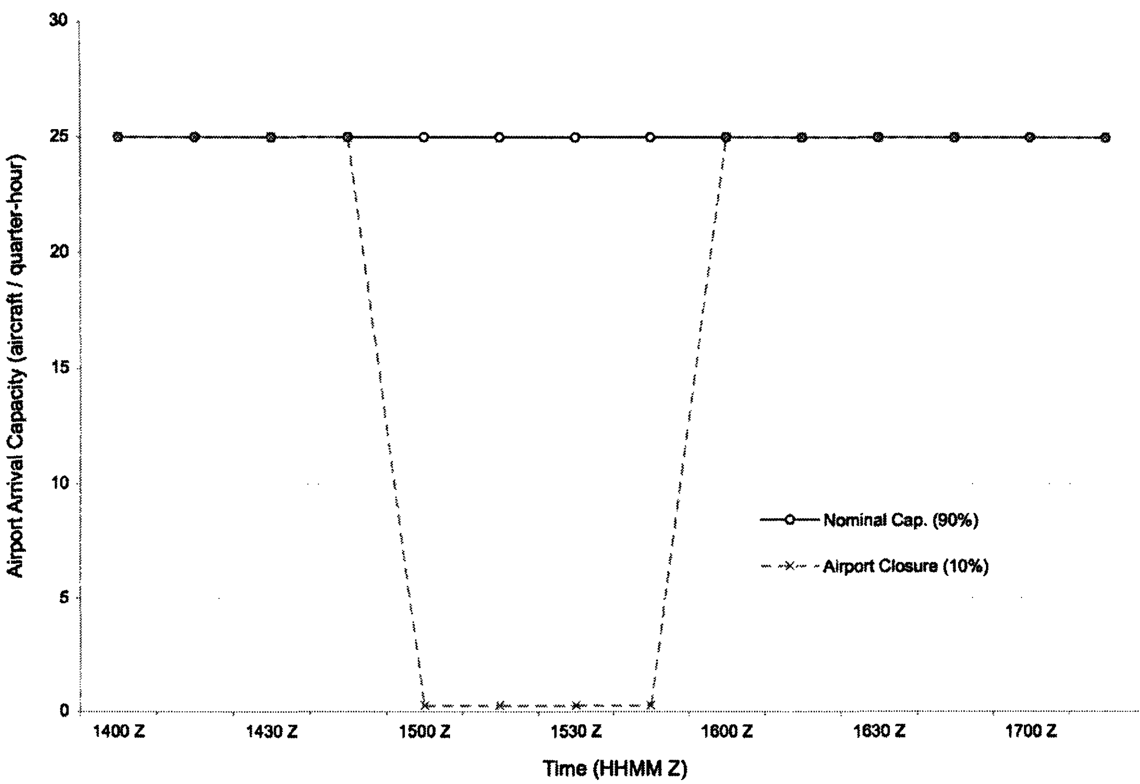

[3] Terrab and Odoni (1993)

Figure 2-1 shows an example of a simple scenario with two profiles. Under Profile #1, the airport maintains a nominal capacity, while Profile #2 represents the passing of a hypothetical storm, for which the PAAR falls to 0 arrivals/period at 1500 Z and then increases to 25 after 1600 Z. At the time this scenario is forecast (1300 Z), the likelihood of Profile #1 is 90% and that for Profile #2 is 10%.

Over time, new information may become available which makes a revision of the arrival capacity forecast possible. Richetta and Odoni (1994) and Mukherjee (2004) model the set of forecasts that may occur over time as a *scenario tree*. The branches of the tree represent individual capacity scenarios and nodes indicate an update, or revision, to the arrival capacity forecast, which results in a new scenario. From a given point on a branch of the tree, the information available regarding the airport arrival capacity will change in a manner defined by one of the continuous paths leading from that point to a leaf on the right.

For the simple example discussed above, it is assumed that the arrival capacity forecast could be revised by direct observation at 1500 Z – one could simply check the actual arrival rate at that time. The scenario tree for this hypothetical example is shown in Figure 2-2 (page 28). The initial arrival capacity scenario, which contains two profiles, is represented as the branch on the left side and the node at 1500 Z signifies the availability of a new forecast. This revised forecast results in one of two new scenarios; the upper branch indicates that Profile #1 will occur and the lower branch Profile #2. The likelihoods of the revised forecast resulting in each of these new scenarios are 90% and 10%, respectively. On the right side of the tree, each leaf represents an arrival capacity scenario that contains a single profile and corresponds to the realization of that profile.

Figure 2-2 An example of a scenario tree

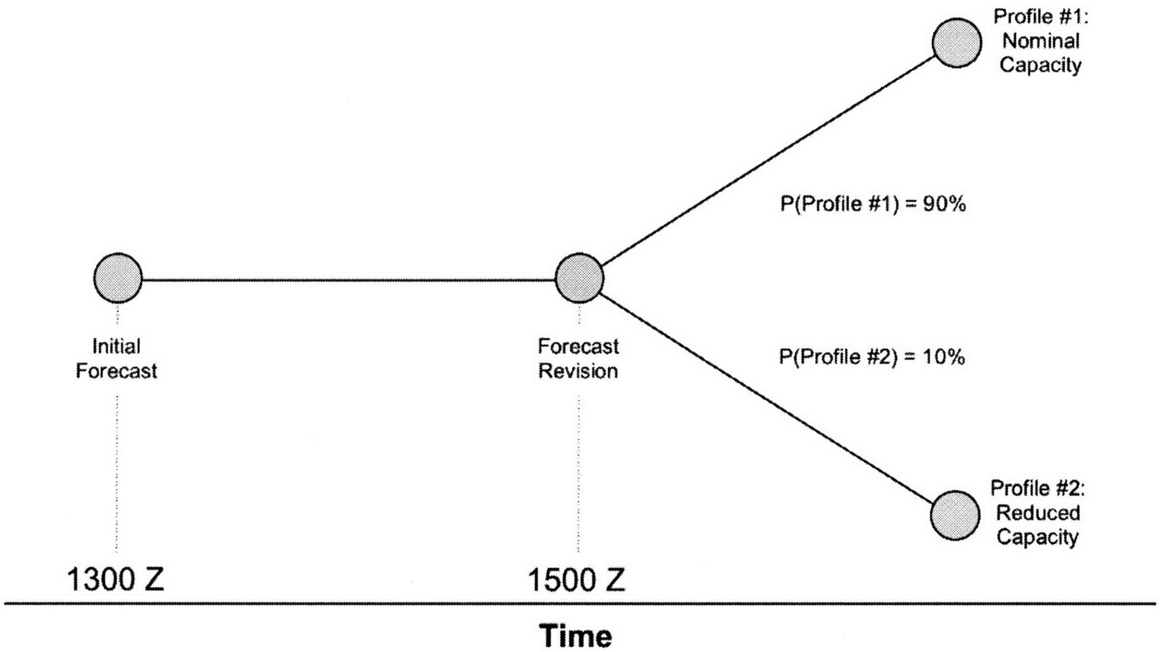

Arrival Capacity Forecasts as Decision Trees

The nodes of a scenario tree could be alternately described as the chance nodes of a decision tree, where a forecast will follow a new branch at each node with a probability defined by the relative likelihoods of the respective capacity profiles. The use of scenario trees to model dynamic arrival capacity forecasts assumes that a revision to the arrival capacity forecast will only occur at the nodes of a tree. This assumption allows for an important simplification in the modeling of a GDP. If the arrival capacity forecast does not change between the nodes of the tree, then decisions regarding a GDP are best made at the nodes, when the capacity forecast has just been revised but there have not been any additional flight departures. Thus, with this assumption, the design of a GDP can be represented as a set of sequential decisions defined by the nodes of a scenario tree, such as the one shown in Figure 2-3.

Figure 2-3 The scenario tree as a decision tree

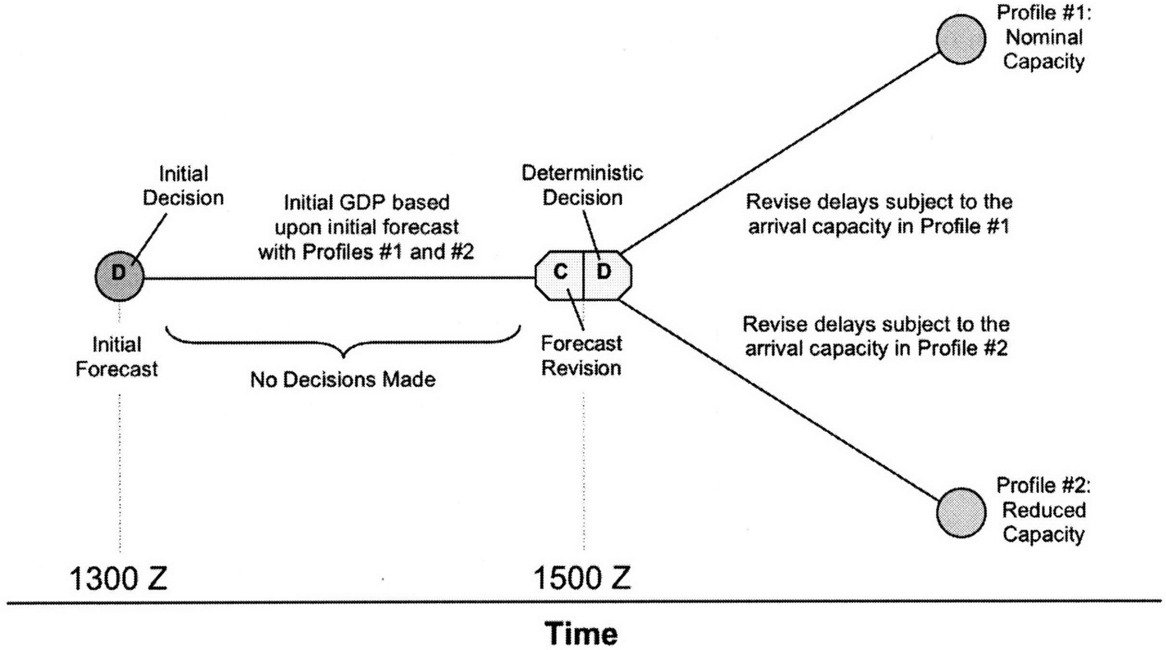

The decision tree framework carries an important implication regarding how a GDP is modeled. GDPs are dynamic decision processes in which a program might be initiated and then revised one or more times as new information becomes available. Each decision to initiate, revise, or cancel the GDP – or even to take no action at all – will impact the alternatives available at later times. Reducing the design of a GDP to a problem of making choices at the nodes of a scenario tree allows for the explicit and simultaneous consideration of the initial GDP, as well as future opportunities to revise the assigned ground delays. And, as shown by the development of various models in the SAGHP literature, it is possible to use optimization techniques to design an optimal GDP subject to this dynamic decision framework.

The Development of Scenario Trees

Additional research has sought to identify methods of developing arrival capacity scenario trees for use in practice. MIT Lincoln Laboratory (2004) discusses the development of a stochastic forecast for the clearing time of marine stratus at San Francisco International Airport (SFO). In the presence of low cloud ceilings, flight arrivals at SFO are restricted to a single runway, which may halve the arrival capacity of the airport. Liu et al. (2006) propose a more generalizable clustering algorithm that develops an arrival capacity scenario tree based upon the historical acceptance rates of an airport. Although the Liu et al. model is applicable to a greater number of airports, it may not yield predictive forecasts that are sufficient for use with optimization methods (Liu 2007). The development of additional, practical methods to construct scenario capacity trees is an area of ongoing research.

Section 2.3.2 SAGHP Optimization Models

In the academic literature, a series of models have been proposed that use optimization methods to solve the SAGHP subject to a stochastic and/or uncertain forecast of the airport arrival capacity, which may be given as a scenario tree. The various models differ in the assumptions each makes in regards to the arrival capacity forecast. Deterministic models only consider forecasts with a single profile and stochastic and static models are limited to forecasts with a single arrival capacity scenario. The most advanced models in terms of arrival capacity consider both a stochastic and dynamic forecast. However, as compared to the deterministic models, the static and stochastic models are more limited in regards to the types of delay cost functions they can consider and the type of optimization method they use to solve for the assignment of ground delays.

As mentioned previously, Terrab and Odoni (1993) propose an IP formulation for the SAGHP for a deterministic and static forecast of the airport arrival capacity. Their model divides time into a set of discrete and continuous periods and is compatible with a scenario that contains a single profile. Terrab and Odoni also propose a dynamic programming (DP) model for a stochastic and static forecast in which the arrival capacities are represented by a set of profiles, or a capacity scenario. For both models, the total delay cost of a solution is the expected value of the sum of ground and airborne delay costs for each flight[4]. Delay cost is assumed to be a non-linear function of delay time that may also vary by aircraft.

Richetta and Odoni (1993) propose a stochastic and static IP formulation for the SAGHP. This formulation is more tractable than Terrab and Odoni for larger problems. However, it assumes that the cost function for airborne delay is both linear and identical for all aircraft. This assumption, which is repeated by later SAGHP models, and its ramifications for the solution to the SAGHP will be discussed further in the next section.

Ball et al. (2003) also propose a stochastic and static SAGHP model with linear costs that has the additional advantage of yielding an integer solution when solved as a linear program (LP). Kotnyek and Richetta (2006) show that the Ball et al. formulation is a special case of the Richetta-Odoni model and that Richetta-Odoni also solves to an integer solution as an LP if the cost of ground delay to a flight is non-decreasing in time. Integer solutions are important for GHP models because aircraft, for practical purposes, must be treated as whole entities. By solving the SAGHP as an LP, the Ball et al. and

[4] For the deterministic model, there is no airborne delay

Richetta-Odoni models solve as or more efficiently than other conceivable models that might use IP, MIP (mixed-integer program), or, especially, DP formulations.

Other SAGHP models consider dynamic forecasts of arrival capacity that may change over time, such as one represented by a scenario tree. Richetta and Odoni (1994) propose a static and partially-dynamic IP model. This model is able to revise the ground delay assigned to aircraft if the arrival capacity forecast is revised during the course of a GDP. However, a limitation of the model is that it is unable to revise ground hold strategies by releasing aircraft previously held on the ground if the forecast of arrival capacity is revised upward. Mukherjee (2004) proposes a stochastic and fully-dynamic MIP model that overcomes these limitations. Mukherjee also demonstrates that by including a super-linear term in the cost of ground delay in the objective function, it is possible to favor solutions that feature distributions of ground delay that might be more equitable. As noted by Odoni (1987) and Vossen (2002), equitable, or fair, allocations of ground delay may be preferable in practice.

Several models have also been proposed in the literature that extend this formulation. Mukherjee (2004) proposes a second stochastic/dynamic MIP formulation that also considers the capacity of the airport arrival fixes. Finally, Mukherjee and Hansen (2007) propose a stochastic/dynamic MIP formulation that considers the capacity of the airborne arrival queue. However, this last model was published after the analyses for this thesis were performed and is not included in the comparisons performed in Chapters 3, 4, and 5.

Section 2.3.3 Models of Delay Cost

Figure 2-4 (page 34) summarizes the various optimization models in the SAGHP literature. Although the Terrab and Odoni models are the least sophisticated in regards to the modeling of arrival capacity, they are the most general in terms of the types of cost functions they consider. On the other hand, the formulations proposed by Mukherjee are advantageous because they consider dynamic strategies in which ground delays are revised as new information becomes available. In comparison to the static and partially dynamic SAGHP models, Mukherjee identifies solutions with a lower total expected cost[5]. However, this requires that the cost of delay – and the objectives of a GDP – be expressed in a form that is compatible with the Mukherjee model. For the design of GDPs in practice, the Mukherjee models, as well as other dynamic SAGHP models in the literature, have four general limitations.

First, the dynamic SAGHP models, such as Mukherjee (2004) and Richetta-Odoni (1994), are limited to delay costs that can be expressed as a linear or a non-linear function of the ground delay plus a linear function of the airborne delay experienced by each aircraft. This restriction follows an initial set of assumptions made by Richetta and Odoni (1993) that model the marginal cost of airborne delay per aircraft as a constant value. However, in practice, delay costs may be much more variable. For example, airborne delays of short duration may have a relatively small cost to NAS users, while longer delays may result in costly flight diversions. Furthermore, some costs may be a function of the total time that a flight is delayed, such as those due to missed connections of passengers, crew, and equipment.

---

[5] Mukherjee (2004)

Figure 2-4 A Summary of SAGHP Models

Model	Type	Arrival Capacity Forecast	Capacity-Constrained Elements	Delay Costs[4]
Terrab and Odoni 1993[1] (1)	IP	Deterministic, Static	Airport Arrivals	Ground: Non-linear[5] Airborne: n/a
Terrab and Odoni 1993[1] (2)	DP	Stochastic, Static	Airport Arrivals	Ground: Non-linear[5] Airborne: Non-linear[5]
Richetta and Odoni 1993	LP[2]	Stochastic, Static	Airport Arrivals	Ground: Non-Linear[5] Airborne: Linear
Ball et al. 2003	LP	Stochastic, Static	Airport Arrivals	Ground: Linear Airborne: Linear
Richetta and Odoni 1994	MIP	Stochastic, Dynamic[3]	Airport Arrivals	Ground: Non-Linear Airborne: Linear
Mukherjee 2004[1] (1)	MIP	Stochastic, Dynamic	Airport Arrivals	Ground: Non-Linear[6] Airborne: Linear
Mukherjee 2004[1] (2)	MIP	Stochastic, Dynamic	Airport Arrivals and Airport Arrival Fixes	Ground: Non-Linear[6] Airborne: Linear
Mukherjee and Hansen 2007	MIP	Stochastic, Dynamic	Airport Arrivals and the Airport Arrival Queue	Ground: Non-Linear[6] Airborne: Linear
Hanowsky 2007	MIP	Stochastic, Dynamic	Airport Arrivals and the Airport Arrival Queue	Ground: Non-Linear Airborne: Non-Linear

[1] Terrab and Odoni (1993) and Mukherjee (2004) each propose multiple SAGHP models
[2] Richetta and Odoni (1993) solves as an LP for non-decreasing marginal ground delay cost
[3] Richetta and Odoni (1994) is partially-dynamic
[4] Delay costs refer to the cost of a delay as a function of the duration of delay experienced by each flight
[5] Terrab and Odoni (1993) and Richetta and Odoni (1994) consider functions of delay cost that may vary by flight. In the other models, each flight is identical from the perspective of how costs are accrued.
[6] Mukerjee (2004) and Mukherjee and Hansen (2007) may not solve efficiently as an IP for all non-linear ground delay cost functions

Second, the cost of delays may also vary by aircraft. The stochastic/dynamic SAGHP models in the literature require that all aircraft be equal from a cost perspective. However, in practice, costs may not be experienced equally, as larger aircraft require more fuel and crew members and also carry more passengers; long distance flights may be more susceptible to diverting as a result of airborne delays; and flights with connecting passengers may be more sensitive to delays, in general, than other flights. It should also be noted that the initial description of the GHP by Odoni (1987) and the model proposed by Richetta and Odoni (1993) assume that delay costs may vary by aircraft.

Third, the solutions to models that only seek to minimize the total expected delay cost, such as those proposed by Mukherjee, permit arrival queue lengths that might be unacceptable to the traffic manager in practice. These models assign delay to minimize a weighted average of the flight delay costs under each profile. For a profile that represents a significant reduction in arrival capacity but for which the likelihood of being realized is relatively low, the weighted cost may not be significant enough to induce the assignment of ground delay. However, in practice, it may be preferable to the traffic manager to avoid the possibility of extensive airborne arrival queues, even if this would increase the total expected delay cost.

Fourth, the Mukherjee model, specifically, may not solve efficiently as a MIP if ground delay or costs are non-linear. Although Mukherjee demonstrates that his formulation may consider some non-linear cost functions, adaptation of the model to consider arbitrary non-linear ground delay cost functions is not trivial. By experimentation, it was found that certain non-linear functions of ground delay cost may cause the model to fail to resolve to an integer solution in a timely manner. Thus, while the Mukherjee formulation is the most advanced SAGHP model in its consideration of stochastic and dynamic arrival capacity forecasts, it is also the most restrictive in terms of the types of delays costs that may be considered. The model proposed in this thesis solves more efficiently the stochastic/dynamic SAGHP for a wider range of cost functions than Mukherjee and other prior models.

Section 2.4 Conclusion

The design of GDPs in practice is subject to stochastic and dynamic forecasts of the airport arrival capacity. A series of models have been proposed in the academic literature that solve for the optimal assignment of ground delay given such a forecast. However, these models have certain gaps or deficiencies that may act as major barriers to their use in practice. In this thesis, a new SAGHP model will be proposed that considers both a stochastic and dynamic arrival capacity forecast and a more generalized model of delay cost.

Chapter 3 A Mixed-Integer Programming Model for the Stochastic and Dynamic Single-Airport Ground Holding Problem with Non-Linear Cost Functions

During periods of inclement weather, the FAA may implement a Ground Delay Program, or GDP, to manage the rate of flight arrivals at a capacity-constrained airport. The objective of a GDP is to avoid extensive airborne delay by proactively delaying flights on the ground before they depart. Ground delay is preferable to airborne delay because delays on the ground are less costly, along several dimensions, than those in the air. In order to assign ground delays, however, a GDP must be initiated before flights depart, at which time the future capacity of the destination airport may be uncertain.

A series of models have been proposed in the academic literature to design a GDP subject to uncertain capacity. Two important criticisms of these models are that they permit very large airborne queues, which may be unacceptable to air traffic managers in practice, and that they are unable to consider non-linear airborne delay costs, which are more representative of how costs are accrued by users of the NAS. This chapter proposes a new model to dynamically assign ground delay to flights that are scheduled to arrive at an airport with uncertain arrival capacity, subject to constraints on size of the airborne queue and with non-linear flight delay costs. The model is applied to a hypothetical example and is shown to be more general and more applicable than those in the existing literature.

This chapter is presented in four sections: §3.1 identifies the advantages offered by the proposed model over those in the literature; §3.2 proposes and discusses the model formulation; §3.3 describes how the model might be applied in practice and, for two simple examples, compares its solution with those of other models; and §3.4 summarizes the contribution of the proposed model to both the literature and the design of GDPs in practice.

Section 3.1 Contributions of the Extended Cost Model

The model presented in this chapter extends a body of literature on the single-airport ground holding problem (SAGHP)[6]. Existing models in the literature, most notably Richetta and Odoni (1994) and Mukherjee (2005), identify an optimal assignment of ground delay subject to a stochastic and dynamic forecast of the airport arrival capacity, which is provided in the form of a scenario tree. However, a limitation of the existing models is that the types of objective functions they permit are unable to capture several important considerations of the design of a GDP in practice, such as costs that occur due to extended airborne delays or the size of the airborne arrival queue. The model presented here improves upon previous stochastic and dynamic SAGHP models and offers four features that make it attractive for use in the design of GDPs in practice.

The first is that the model considers a wider range of objective functions than previous models in the literature. Whereas existing models are limited to an objective of the sum of a linear or non-linear function of ground delay time and a linear function of airborne delay time, this new model minimizes the sum of linear or non-linear functions of the airborne, ground, and cumulative delay time of each flight. As this model extends

[6] §2.3

the range of delay cost function that may be considered in the design of a GDP it is, therefore, referred to as the "extended cost model," or ECM.

A second feature of the ECM is that it is capable of distributing all delay – not just ground delay – more evenly among aircraft. Given two solutions with equal total delay cost, the solution that distributes delay more evenly may be preferable in practice. Prior models have shown that it is possible to identify solutions that distribute ground delay more evenly. However, these models are unable to consider the distribution of airborne delay, which is relevant because flights that are delayed during a GDP may experience both ground and airborne delay. The ECM is able to identify solutions that distribute ground, airborne, or the total amount of delay assigned to each flight more evenly.

A third feature of the ECM is that it can be solved in a reasonable amount of time for use in practice. In nearly all examples tested, including those for a variety of different non-linear delay cost functions, the LP relaxation of the ECM yields an integer solution. In those cases for which an integer solution is not found for the LP relaxation, an exact, optimal solution is quickly identified using a branch-and-bound algorithm. For the examples discussed in this thesis, model run time is often on the order of a few seconds and the largest problem solved requires less than ten minutes to find an optimal solution.

A fourth feature of the ECM is that it is able to avoid solutions that could result in excessive airborne delays and long arrival queues. There may exist solutions that minimize the total expected cost but allow for the possibility of lengthy arrival queues that might not be acceptable to traffic managers in practice. The ECM allows the traffic manager to place limits on the size of the arrival queue that is permitted in the optimal solution.

Section 3.1.1 A Model of Delay Cost

The model of delay cost that is used by the ECM takes the perspective a flight that experiences delay during a GDP. The cost of the delay to the flight may be a consequence of one or more of many different possibilities (Figure 3-1), such as the fuel consumed during airborne delay. Another example is labor, which incurs cost as a function of the cumulative delay time. *Cumulative* delay refers to the total amount of time a flight is delayed and is defined as the sum of air and ground delay times (Figure 3-2).

Figure 3-1 A classification of various delay costs experienced by NAS users

Cost	Description	Delay Type	Cost Function
Cancellations	Cost of flight cancellation due to excessive ground delay	Ground	Non-Linear
Diversions	Cost of diverting to an auxiliary airport due to excessive airborne delay	Airborne	Non-Linear
Fuel	Fuel costs due to extra flight time	Airborne	Linear
Labor	Labor costs due to delayed arrival	Cumulative	Linear or Non-Linear
Missed Connections	Costs to reroute passengers, crew, and equipment that missed scheduled connections	Cumulative	Non-Linear

Three types of cost (in addition to fuel and labor) shown in Figure 3-1 are of interest because each is non-linear with respect to a different type of delay time: the likelihood of a cancellation may increase with the duration of ground delay, of diversions with airborne delay, and of missed connections (of passengers, crew, and equipment) with cumulative delay. For example, consider two possible allocations of airborne delay, one that delays one flight for 45 minutes and the second that delays three flights for 15 minutes each. Although the total delay time in each allocation is the same, the lengthy airborne delay in the first might result in a costly diversion that would not occur in the second.

A key assumption made by this model of delay cost is that the three categories of cost, ground, airborne, and cumulative, are mutually exclusive and collectively exhaustive. Each of the various costs incurred by a delayed aircraft must be assigned to exactly one category. For example, labor cost is assigned to cumulative delay and not to either the airborne or ground delay cost categories. In order to prevent a double-counting of such costs by the ECM, it is important that costs be allocated into categories properly before the model is applied.

Figure 3-2 The relationship between ground, airborne, and cumulative delay

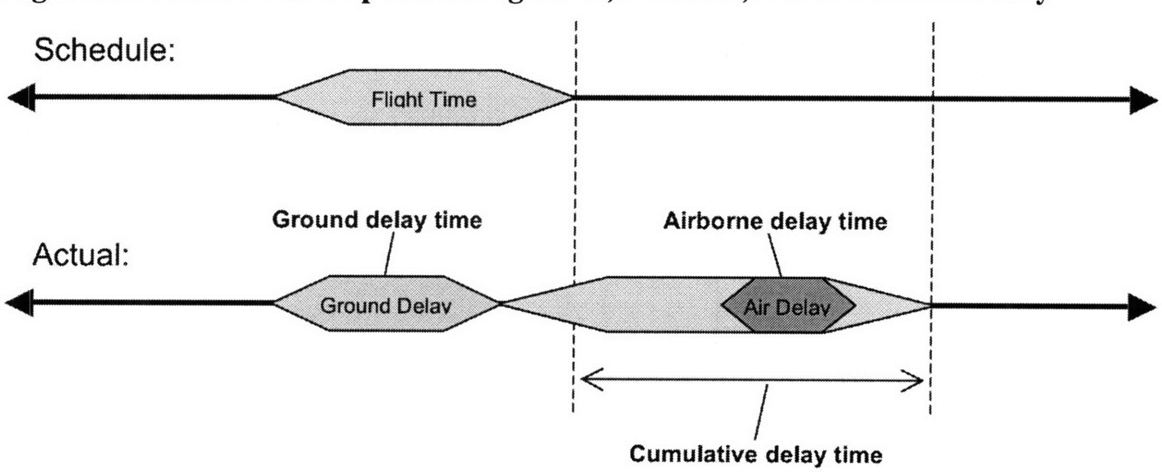

Section 3.1.2 The Distribution of Delays

A second, more subtle benefit of the use of non-linear cost functions is to identify solutions that distribute delay more evenly across aircraft participating in a GDP. For a given objective function, there may exist multiple optimal solutions with the same total delay cost. Consider the previous example with two possible solutions: the first that delays one aircraft for three periods of 15 minutes and the second that delays three aircraft for one period each; except, in this case, let the delays be taken on the ground. Assuming that ground delay has a cost of one unit per period, these two solutions would have an identical cost. However, from the perspective of the traffic manager, the latter solution may be preferred because it assigns some delay to each aircraft rather than all delay to a single aircraft.

Mukherjee (2005) demonstrates that if there exist multiple optimal solutions, then it is possible to favor those solutions that assign ground delay more evenly across flights by including a super-linear term in the objective function. However, it may be preferable to consider the distribution of cumulative delay, rather than ground delay, because it is possible for a flight to receive both air and ground delay. The ECM is able to consider cost functions that incorporate a super-linear term into the cost of cumulative delay, thus distributing both ground and airborne delay more evenly.

Section 3.1.3 An Airborne Arrival Queue Capacity

In practice, traffic managers may assess alternatives for strategic ATFM decisions by placing limits on airspace usage, such as a cap on the number of aircraft simultaneously permitted in a sector or volume of airspace. These limits may differ across sectors and may vary over time for the same sector, as inclement weather can reduce the amount of traffic a volume of airspace may accommodate. For the design of GDPs subject to an uncertain airport arrival rate capacity, planning for a specific capacity rate might lead to situations in which there is a low, but appreciable risk of exceeding the maximum number of aircraft allowed to be in an airborne queue at the same time. The SAGHP models in the literature, which seek only to minimize the expected delay cost, permit solutions in which the size of the airborne arrival queue might exceed acceptable limits. To avoid these solutions, the ECM allows the traffic manager to specify a capacity, or maximum limit, for the arrival queue under each profile.

Section 3.2 The Extended Cost Model MIP Formulation

Five key concepts for the design of a GDP are described in §3.1:

- Dynamic decision-making
- Uncertain arrival capacity
- Prevention of unacceptable airborne queue sizes
- Minimization of the expected cost of delays to NAS users
- The distribution of delay among aircraft

This section presents the formulation of a mixed-integer mathematical program for the design of a GDP for a dynamic and stochastic arrival capacity forecast, with non-linear cost functions, and subject to constraints on the length of the airborne arrival queue. For convenience, the ECM is repeated without the explanatory text in the appendix.

Section 3.2.1 Notation

Consider an airport for which there is a list of scheduled flight arrivals and a dynamic and stochastic forecast of the arrival capacity. The list of scheduled arrivals contains a set of flights F that are scheduled to arrive at the airport. Each flight $f \in F$ has a scheduled departure time and a scheduled arrival time. The arrival capacity forecast is provided as a scenario tree with a set of capacity profiles Q and a set of capacity scenarios B. Each capacity profile $q \in Q$ represents a profile of PAARs over time and has a likelihood of being realized given by probability p_q, with $\sum p_q = 1$.

Let time be divided into a series of consecutive time periods $\tau = \{0,\ldots,T\}$, where each period $t \in \tau$ represents a distinct period of time. The initial period $t = 0$ is the time at which the current decision regarding a possible GDP is being made and the final period of time $t = T$ is assumed to be far enough in the future as to accommodate a GDP of any length. In practice, GDPs are limited to a single day of demand, so a reasonable time horizon would be between 8 and 12 hours. By convention, time periods have equal duration (for example, 15 minutes), although the ECM can be applied to any set of distinct and sequential time periods.

Arrival demand and capacity is defined by the set of time periods. Let the scheduled departure and arrival times of flight f be $SDT_f \in \tau$ and $SAT_f \in \tau$. For convenience, the scheduled flight time en route STE_f is defined as the difference between these times:

$$STE_f = SAT_f - SDT_f \quad \forall \ f \in F$$

The planned airport arrival capacity under profile q during time period t is M_{qt}, which is assumed to be both non-negative and integer. It is assumed that the capacity of the final time period $t = T$ is sufficient to handle all arrivals.

$$M_{qT} = |F| \quad \forall \; q \in Q$$

As part of the arrival capacity scenario tree, the forecast will be updated at specific times, as represented by the nodes of the scenario tree. Each scenario $b \in B$ is associated with a start time period $b_s \in \tau$ and an end time period $b_e \in \tau$, which correspond to the start and end nodes of the scenario. For the example shown in Figure 3-3, the start time period of b_1 corresponds to 1300 Z and the end time period to 1500 Z.

Figure 3-3 Arrival capacity scenario tree notation

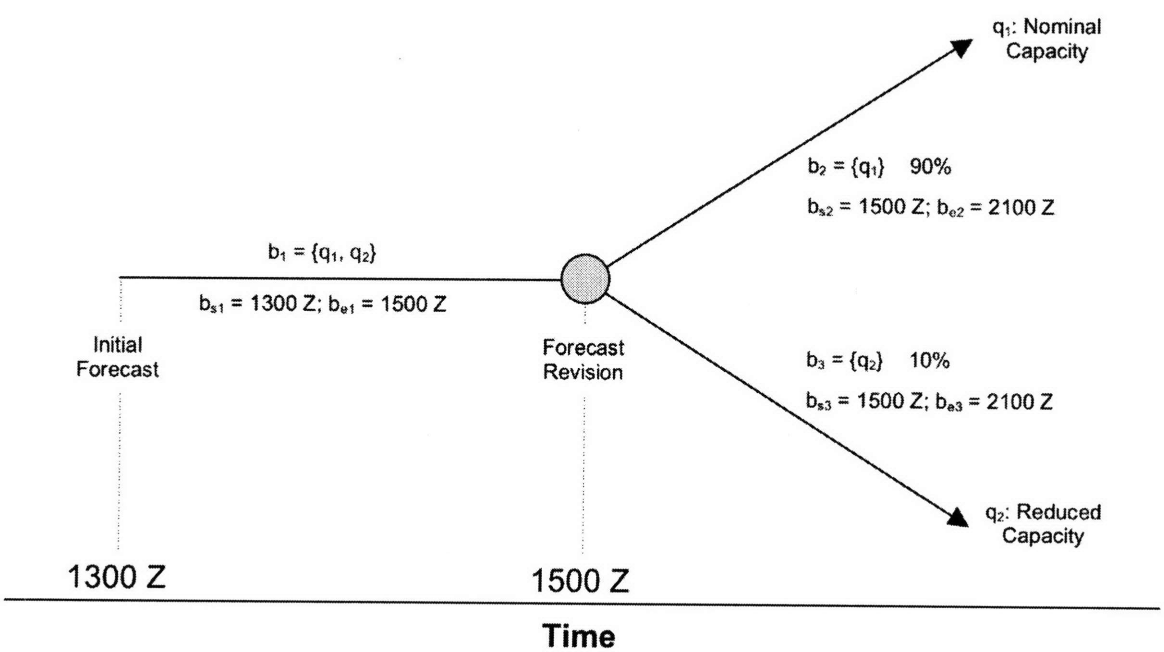

Each scenario b also contains the set of profiles that could be realized following b, such that $b = \{q_i \ldots q_j\} \subseteq Q$. Thus, each profile q is an element of the set of all profiles Q and of one or more scenarios as defined by the scenario tree. For example, in Figure 3-3, profile q_1 is an element of scenarios b_1 and b_2.

The Cost of Delay

Consider a flight *f*, which may experience ground or airborne delay or both before it arrives at an airport. The total cost of delay to *f* is defined as the sum of the costs that are incurred as a result of the ground delay, airborne delay, and cumulative delay experienced by *f*. For each of these three types of delay, cost will vary by the duration of delay and may be expressed as the sum of incremental amounts. For example, the cost of n periods of airborne delay for flight *f* is $\sum_{i=1}^{n} \alpha_i$, where α_i is the incremental cost of the i^{th} period of airborne delay and $i \in \tau$. Similarly, γ_i is the incremental cost for the i^{th} period of ground delay and χ_i is the incremental cost for the i^{th} period of cumulative delay. Thus, the total delay cost to a flight that receives m periods of ground delay and n periods of airborne delay is $\sum_{1}^{m} \gamma_i + \sum_{1}^{n} \alpha_i + \sum_{1}^{m+n} \chi_i$.

It is assumed that the cost when no delay is incurred is 0 units and that the incremental costs are non-decreasing:

$$\alpha_0 = \gamma_0 = \chi_0 = 0$$

$$\chi_t \geq \chi_{t-1} \qquad \forall \ t \geq 1$$

$$\alpha_t \geq \alpha_{t-1} \qquad \forall \ t \geq 1$$

$$\gamma_t \geq \gamma_{t-1} \qquad \forall \ t \geq 1$$

This assumption is a reasonable approximation to the accrual of delay costs to NAS users, where short delays can often be easily absorbed by "buffer times" in schedules, but long delays have a more significant impact.

Section 3.2.2 Decision Variables

The decision variables for the ECM indicate the departure time period, arrival demand time period, and landing time period of each flight, where the arrival demand period is the time at which a flight would land if not otherwise delayed in the air. The variables are d_{ftq}, δ_{ftq}, and λ_{ftq}, which represent the fraction of flight f that departs, demands arrival, or lands, respectively, by time t and under profile q; only d_{ftq} is required to be binary.

$$d_{ftq} \in \{0,1\} \qquad \forall \ f \in F; \ t \in \tau; \ q \in Q$$

$$\delta_{ftq} \in [0,1] \qquad \forall \ f \in F; \ t \in \tau; \ q \in Q$$

$$\lambda_{ftq} \in [0,1] \qquad \forall \ f \in F; \ t \in \tau; \ q \in Q$$

Although included as a decision variable in the ECM, the flight arrival demand variable δ_{ftq} is fully determined by the departure variable d_{ftq} and the scheduled time en route STE_f, which the ECM assumes to be deterministic. Although the ECM could be reformulated without δ_{ftq}, it was noted during initial experimentation that retaining the arrival demand variables does not significantly increase the run time of the model. Furthermore, including these variables facilitates the future adaptation of the ECM to consider stochastic en route times.

Figure 3-4 Sample decision variables for one flight

Time Period	0	1	2	3	4	5	6	7	8	9	10
d_{ftq} = Departure	0	0	1	1	1	1	1	1	1	1	1
∂_{ftq} = Arrival Demand	0	0	0	0	0	0	0	1	1	1	1
λ_{ftq} = Landing	0	0	0	0	0	0	0	0	0	1	1

Ground delay time **Time en route** **Air delay time**

Figure 3-4 shows a possible solution to the ECM corresponding to a flight f with $SDT_f = 1$ and $SAT_f = 6$ and a capacity profile q. The first row of values indicates the departure status of f over time assigned by the solution. The value of the departure variable for the initial time period $t = 0$ is zero, which indicates that f has not departed prior to the initiation of the program. The change in the value of the indicator from zero to one at $t = 2$ indicates that flight f departs during the second time period in this solution. The difference between the scheduled departure time and the departure time assigned by the model, which is one period, is ground delay.

The second and third rows in the figure correspond to the arrival demand and landing variables. As shown, flight f would have been able to land at time $t = 7$, but does not actually land until time $t = 9$. The difference between the arrival demand and landing times assigned by the model, which is two periods, is airborne delay. Lastly, the difference between the assigned landing time and the scheduled departure time, which is three periods, is cumulative delay.

The departure, arrival demand, and landing indicator variables have three properties:

1. Their values are non-decreasing
2. There is a hierarchy across variables: a flight must depart before it demands arrival, and it must demand arrival before it lands
3. Values must be binary for practical reasons

The third property is of special importance. Fractional values for the departure indicator variables are impractical for use in the design of a GDP. For example, if $d_{ftq} = ½$ in the optimal solution, it would indicates that one-half of flight f should depart by time t. As discussed in §2.3.3, the branch-and-bound algorithm may not resolve these fractional values efficiently. Thus, requiring that variables take binary values is necessary for practical reasons, but also requires that the ECM be solved as an MIP and not an LP.

The arrival demand and landing variables are not required to be binary in the ECM. This assumption is consistent with the use of GDPs in practice, where only flight departure times are assigned. However, this also means that the arrival demand and landing variables could take fractional values in the optimal solution. Should this occur, the departure times assigned by the model would still be optimal if and only if there exists an alternate solution with the same assignment of departure times and an identical objective function value, but with binary values for the arrival demand and landing variables.

Claim: If there exists an optimal solution to the ECM with binary values for d_{ftq}, *then there exists an optimal solution with binary values for* δ_{ftq} *and* λ_{ftq}.

Corollary: If there exists an optimal solution to the ECM and if the incremental flight delay cost parameters adhere to the following relationship:

$$\alpha_i + \chi_n \le \alpha_j + \chi_j \qquad \forall \quad n;\ i < j$$

then there exists an optimal solution in which flight arrivals are first-come-first serve (FCFS).

The claim and corollary indicate that the solution identified by the ECM is optimal if the flight arrival process is FCFS, even though the model does not require it. Proofs of both are available in the appendix.

Auxiliary Decision Variables

An additional set of decision variables is used in the ECM to calculate the airborne delay cost of each flight. Let

$$Z_{fstq} \in \Re \qquad \forall\ f \in F;\ s \in 0 \ldots T - SAT_f - 1;\ t \in s \ldots T - SAT_f - 1;\ q \in Q$$

These variables will be discussed further in the next section.

Section 3.2.3 Objective Function

The objective function of the ECM minimizes the total expected delay cost:

$$\textit{Minimize} \qquad \sum_Q \left[p_q \times \sum_F \left(CDC_{fq} + ADC_{fq} + GDC_{fq} \right) \right] \tag{3.01}$$

Where CDC_{fq}, ADC_{fq}, and GDC_{fq} represent the cumulative, air, and ground delay costs, respectively, for flight f under profile q and are defined in equations (3.02) – (3.06) below.

$$CDC_{fq} = \sum_{t=SAT_f}^{T-1}\left(\left(1-\lambda_{ftq}\right)\times\chi_{t-SAT_f+1}\right) \qquad \forall\ f\in F;\ q\in Q \qquad (3.02)$$

$$GDC_{fq} = \sum_{t=SDT_f}^{T+SDT_f-SAT_f-1}\left(\left(1-d_{ftq}\right)\times\gamma_{t-SDT_f+1}\right) \qquad \forall\ f\in F;\ q\in Q \qquad (3.03)$$

$$ADC_{fq} = \sum_{t=SAT_f}^{T-1}\left(\left(1-\lambda_{ftq}\right)\times\alpha_{t-SAT_f+1}\right) - \sum_{s=0}^{T-SAT_f-1}\sum_{t=s}^{T-SAT_f-1} Z_{fstq} \qquad \forall\ f\in F;\ q\in Q \qquad (3.04)$$

$$Z_{fstq} = \left(\alpha_{t-s+1}-\alpha_{t-s}\right)\times\min\left[1-\lambda_{f(t+SAT_f)q},\ 1-d_{f(s+SDT_f)q}\right] \qquad (3.05)$$

$$\forall\ f\in F;\ s\in 0\ldots T-SAT_f-1;\ t\in s\ldots T-SAT_f-1;\ q\in Q$$

For the purposes of the formulation, equations (3.05) are replaced with constraints (3.06a) and (3.06b):

$$Z_{fstq} \le \left(\alpha_{t-s+1}-\alpha_{t-s}\right)\times\left(1-\lambda_{f(t+SAT_f)q}\right) \qquad (3.06a)$$

$$Z_{fstq} \le \left(\alpha_{t-s+1}-\alpha_{t-s}\right)\times\left(1-d_{f(s+SDT_f)q}\right) \qquad (3.06b)$$

Cumulative delay cost CDC_{fq} and ground delay cost GDC_{fq} are the sum of the costs associated with each incremental period of delay assigned to each flight. For example, consider equations (3.02), which express the cumulative delay cost as the function with two terms. The first term $\left(1-\lambda_{ftq}\right)$ is equal to 1 if and only if flight f does not land by time t for profile q. The second term is the cumulative delay cost increment χ_{t-SAT_f+1}, which is indexed to the scheduled arrival time of f. Equations (3.03) similarly express flight ground delay cost.

Figure 3-5 Delay cost calculation example

Time Period	0	1	2	3	4	5	6	7	8	9	10
Ground Delay		γ_1	γ_2	γ_3	γ_4	γ_5	γ_6	γ_7	γ_8	γ_9	γ_{10}
Cumulative Delay							χ_1	χ_2	χ_3	χ_4	χ_5
Airborne Delay							α_1	α_2	α_3	α_4	α_5
Ground delay		**SDT_f**					**SAT_f**			**Cumulative Delay**	
$1 - d_{ftq}$		**1**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**	**0**
$1 - \lambda_{ftq}$							**1**	**1**	**1**	**0**	**0**

A sample calculation of cumulative and ground delay costs is shown above in Figure 3-5. Recall the previous example of flight decision variables (Figure 3-4, page 48) in which a flight f receives one period of ground delay and two periods of airborne delay. The time periods during which f experiences delay, which are indicated by shading in the figure, correspond to the times for which the $\left(1-d_{ftq}\right)$ and $\left(1-\lambda_{ftq}\right)$ terms equal 1. As the vectors of incremental ground delay and cumulative delay costs are offset by the scheduled departure and arrival times, respectively, the equations (3.02) and (3.03) calculate the costs of cumulative and ground delay as $CDC_{fq} = \chi_1 + \chi_2 + \chi_3$ and $GDC_{fq} = \gamma_1$.

In contrast to cumulative and ground delay costs, there does not exist a convenient time at which to index the airborne delay cost coefficient. Instead, airborne delay cost is expressed in equations (3.04) as the difference of two terms. The first term is similar to (3.02) and (3.03) and represents the cost of air delay cost to flight f that would be incurred if all of the observed delay in arrival time is taken as airborne delay (i.e. if none of the observed delay occurs on the ground). For the flight in the previous example that received two periods of airborne delay and one period of ground delay, the value of the first term is $\alpha_1 + \alpha_2 + \alpha_3$. However, the cost for of two periods of airborne delay is actually $\alpha_1 + \alpha_2$; by itself, the first term overstates the cost of airborne delay.

The purpose of the second term in equations (3.04) is to subtract the delay cost that is overstated by the first term. The sum of the Z_{fstq} decision variables in the second term represents the reduction in delay cost that is due to some of the observed delay in arrival time being taken on the ground and not in the air. Thus, for the given example, the value of the second term is equal to α_3.

Figure 3-6 A sample calculation of the Z_{fstq} values

	$1-\lambda_t$	1	1	1	0	0
$1-d_s$	s \ t	1	2	3	4	5
1	1	α_1	$\alpha_1-\alpha_2$	$\alpha_2-\alpha_3$	$\alpha_3-\alpha_4$	$\alpha_4-\alpha_5$
0	2		α_1	$\alpha_1-\alpha_2$	$\alpha_2-\alpha_3$	$\alpha_3-\alpha_4$
0	3			α_1	$\alpha_1-\alpha_2$	$\alpha_2-\alpha_3$
0	4				α_1	$\alpha_1-\alpha_2$
0	5					α_1

Figure 3-6 shows how the second term calculates the cost reduction. The axes of the table are expressed in periods of time, where s represents the time periods starting with the scheduled departure time (such that $SDT_f = 1$) and t the time periods starting with the scheduled arrival time. The vectors that correspond to the $\left(1-d_{ftq}\right)$ and $\left(1-\lambda_{ftq}\right)$ terms are also indexed by s and t, respectively. The elements of the body of the table represent the difference between consecutive values of α_t and are defined as

$\alpha_{t-s+1} - \alpha_{t-s}$[7]. If s represents the incremental ground delay time and t the incremental airborne delay time, then the value of cell (s,t) is the reduction in the t^{th} period of airborne delay cost due to the s^{th} period of ground delay.

As the second term in equations (3.04) has a negative coefficient, minimizing the objective function (3.01) forces the Z_{fstq} variables to take the maximum possible values permitted by constraints (3.06a) and (3.06b). For non-decreasing airborne delay cost increments, $Z_{fstq} > 0$ only if flight f is assigned at least s periods of ground delay and t periods of airborne delay under profile q. For this example, $Z_{fstq} = 0$ for $t > 3$ (3.06a) or $s > 1$ (3.06b). Thus:

$$Z_{f11q} = \alpha_1$$

$$Z_{f12q} = \alpha_2 - \alpha_1$$

$$Z_{f13q} = \alpha_3 - \alpha_2$$

$$Z_{fstq} = 0 \quad \textit{otherwise}$$

Thus, $\sum Z_{fstq} = \alpha_3$ and $ADC_{fq} = \alpha_1 + \alpha_2$.

Contributions of the Cost Functions

The form of the delay cost functions contributes two key benefits to the study of the GHP. First, the expressions of ground delay cost (3.02) can consider any function that is represented in a piecewise linear form. As discussed in §2.3, existing stochastic and fully-dynamic models in the literature[8] may not consider an arbitrary non-linear

---

[7] Cells for which $s > t$ are excluded from both the table and the summation in the second term because the amount of ground delay assigned to a flight cannot exceed the amount of cumulative delay.
[8] Mukherjee (2004)

ground delay cost function. The ECM is the first model to solve the stochastic and dynamic SAGHP subject to a piece-wise linear ground delay cost function.

A second contribution is the consideration of non-linear cumulative and airborne delay costs. Previous models in the SAGHP literature consider only linear costs. Although equations (3.04) – (3.06) also require a significant number of additional decision variables, experimental results indicate that that the ECM often yields binary solutions when solved as an LP relaxation, which avoids the use of the branch-and-bound algorithm. Non-integer solutions will be discussed in greater detail in Chapter Four and model size is discussed in §3.2.5.

Section 3.2.4 Constraints

The constraints for the ECM are expressed in equations (3.07) – (3.15) below:

$$d_{ftq} \geq d_{f(t-1)q} \qquad \forall\ f \in F;\ t \in 1 \ldots T;\ q \in Q \qquad (3.07)$$

$$\lambda_{ftq} \geq \lambda_{f(t-1)q} \qquad \forall\ f \in F;\ t \in 1 \ldots T;\ q \in Q \qquad (3.08)$$

$$d_{f(t-STE_f)q} = \delta_{ftq} \qquad \forall\ f \in F;\ t \in STE_f \ldots T;\ q \in Q \qquad (3.09)$$

$$\lambda_{ftq} \leq \delta_{ftq} \qquad \forall\ f \in F;\ t \in \tau;\ q \in Q \qquad (3.10)$$

$$\delta_{ftq} = 0 \qquad \forall\ f \in F;\ t \in 1 \ldots SAT_f - 1;\ q \in Q \qquad (3.11)$$

$$\lambda_{fTq} = 1 \qquad \forall\ f \in F;\ q \in Q \qquad (3.12)$$

$$\sum_F \left(\delta_{ftq} - \lambda_{ftq}\right) \leq W_{tq}^{MAX} \qquad \forall\ t \in \tau;\ q \in Q \qquad (3.13)$$

$$\sum_F \left(\lambda_{ftq} - \lambda_{f(t-1)q}\right) \leq M_{tq} \qquad \forall\ t \in \tau;\ q \in Q \qquad (3.14)$$

$$d_{ftq_i} = \ldots = d_{ftq_j} \qquad \forall\ f \in F;\ b \in B;\ q_i, \ldots, q_j \in b;\ t \in \{b_s \ldots b_e\} \qquad (3.15)$$

Recall that the decision variables must be non-decreasing in time and that the departure, arrival demand, and landing of each flight must occur in a specific order. These requirements are reflected by constraints (3.07) – (3.10). Constraints (3.07) and (3.08) ensure that departure and landing indicator variables are non-decreasing in time. Constraints (3.09) set the arrival demand indicator of flight f at time t (in profile q) to be equal to the departure indicator at time $t - STE_f$, effectively requiring that f demands arrival STE_f periods after it departs. Furthermore, as the departure variables d_{ftq} are non-decreasing, the arrival demand variables δ_{ftq} are non-decreasing, as well. Lastly, constraints (3.10) require that each flight lands at or after the time at which it demands arrival.

Constraints (3.11) and (3.12) establish bounds on the assignment of flight departure and landing times. Constraints (3.11) prevent flights from being assigned arrival demand times earlier than scheduled (and, by (3.09), an earlier departure time). Constraints (3.12) require that all flights land by the final time period for all profiles q. As the capacity of the final arrival period for all profiles M_{qT} is equal to the total number of flights, by definition, there exists a solution to the formulation that does not violate (3.12).

Constraints (3.13) and (3.14) require the number of flights demanding arrival and landing to be less than specified limits. Constraints (3.13) specify that the number of aircraft in an airborne arrival queue during time period t in profile q cannot exceed any limits that may be placed on the size of the airborne queue by traffic managers. The values of these limits W_{qt}^{MAX} may be specified for each time period $t \in \tau$ and under each profile $q \in Q$. Constraints (3.14) similarly require that the number of flights that are assigned to land at time t not exceed the planned arrival capacity of the airport specified by the capacity forecast.

Finally, constraints (3.15) are coupling[9], or non-anticipativity, constraints. The formulation contains a set of decision variables for each capacity profile. However, the various profiles also belong to one or more of the branches of the decision tree, such that each branch contains a set of arrival capacity profiles $\{q_i, \ldots, q_j\}$[10]. For the time periods associated with a branch b, coupling constraints require that the assignment of ground delay for all flights must be the same under all of the profiles contained in b.

Section 3.2.5 Model Size

As the ECM may be solved as a MIP, the size of the model, as measured by the numbers of variables and constraints, is of practical consideration; if too large, the model may not solve in a reasonable amount of time, if at all. The ECM requires four types of decision variables, d_{ftq}, δ_{ftq}, λ_{ftq}, and Z_{fstq}. The first three are each defined per flight, time period, and profile, and may be described using mathematical notation as $O(FTQ)$, where F, T, and Q refer to the number of flights, time periods, and profiles. However, the Z_{fstq} variables, which are used for the calculation of the cost of airborne delay, are $O(FT^2Q)$. Similarly, most constraints are also $O(FTQ)$ or less, while those for the calculation of airborne delay cost are $O(FT^2Q)$.

Due to the calculation of airborne delay cost, the numbers of variables and constraints of the ECM increase on the order of the number of time periods squared. As a

[9] Richetta and Odoni (1994)
[10] §3.2.1

result, the size of the model is sensitive to both the total amount of time that is considered and the duration of each time period. It is believed that the size of the examples that are discussed in this dissertation, which require as many as 1 million variables and 2 million constraints, may be representative of the size of some problems that occur in practice[11]. However, further research to revise or reformulate the ECM in order to reduce the size of the model is strongly suggested.

One possibility would be to change how the ECM calculates the cost of airborne delay. As noted previously, the variables and constraints that refer to airborne delay cost (3.06a and 3.06b) increase the size of the model significantly. Although a beneficial feature of the form of these constraints is that the ECM often yields a binary solution when solved as an LP, there may exist an alternate formulation that both solves as an LP relaxation and is of smaller size.

Section 3.3 Demonstration of the Model

To demonstrate the capabilities of the ECM, the following section illustrates its application to two simple, hypothetical examples. In each, the model produces a GDP that minimizes the total expected cost of delays subject to an airborne arrival queue capacity. Furthermore, the ECM is shown to compare favorably to prior models in the literature. It avoids solutions that might be unacceptable in practice and/or results a lower total expected cost.

[11] Details on the implementation are discussed in §3.3.1

Section 3.3.1 Experimental Setup: The Base Case

Experiments using the ECM were run using a base case, which includes a given flight demand list, flight exemptions, arrival queue capacity, delay cost functions, computer platform, and set of models to which the ECM is compared.

Arrival Demand

The airport arrival demand is provided by a list of scheduled flight arrivals. This list contains 348 aircraft that are expected to arrive at Chicago O'Hare International Airport (ORD) between 1500 Z and 1900 Z on June 22, 2005, as forecast by the Enhanced Traffic Management System, or ETMS. ETMS data, which are based on a combination of airline schedules (from the Official Airline Guide, or OAG), filed flight plans, flight-tracking data, and atmospheric information, represent the best objective and comprehensive information available to traffic managers during the design of a GDP.

As the inputs to ETMS are changing over time, the ETMS forecasts of flight arrival times are also subject to change. In this case, the snapshot of forecasted ORD arrivals was taken at 1300 Z on June 22, 2005 and would be representative of the arrival demand forecast available at 1300 Z for a possible GDP starting at 1500 Z. Figure 3-7 (page 60) shows a sample of these data, which includes the departure airport and scheduled departure time of each flight. Flights that are airborne at 1300 Z are indicated by a letter "A" that precedes the departure time.

Figure 3-8 (page 60) illustrates the cumulative arrival demand forecast by time period for the duration of the GDP. Note that the nominal arrival capacity rate of ORD is around 100 aircraft/hour, or 25 aircraft/period. The scheduled arrival demand at 1545 Z and 1715 Z exceeds the nominal capacity; this implies that even without any adverse weather, some delays may occur. Furthermore, as the average demand is 22 aircraft/hour[12], the buffer to accommodate aircraft if the arrival capacity is reduced is very limited.

Figure 3-7 A sample of ten flights from the demand list

#	ACID	TYPE	ORIG	ETD	DEST	ETA
1	53A1186	B735	MSY	A1237	ORD	E1500
2	53A111	B752	BDL	A1236	ORD	E1502
3	10A305	CRJ2	ORF	P1304	ORD	E1502
4	47A126	DC93	MSP	P1403	ORD	E1502
5	53A529	A319	BOS	A1216	ORD	E1504
6	01A829	MD82	BDL	A1230	ORD	E1504
7	01A311	MD83	LGA	A1252	ORD	E1504
8	10A321	CRJ2	RDU	P1319	ORD	E1504
9	01A548	MD82	TUL	P1335	ORD	E1505
10	27A310	B733	AEX	P1303	ORD	E1507

ACID: Aircraft Identification (airline prefix is masked)

TYPE: Aircraft Type

ORIG/DEST: Airport of flight origination/destination

ETD/ETA: Time of departure, (E)stimated, (P)lanned, or (A)ctual

Figure 3-8 Cumulative arrival demand forecast by time period

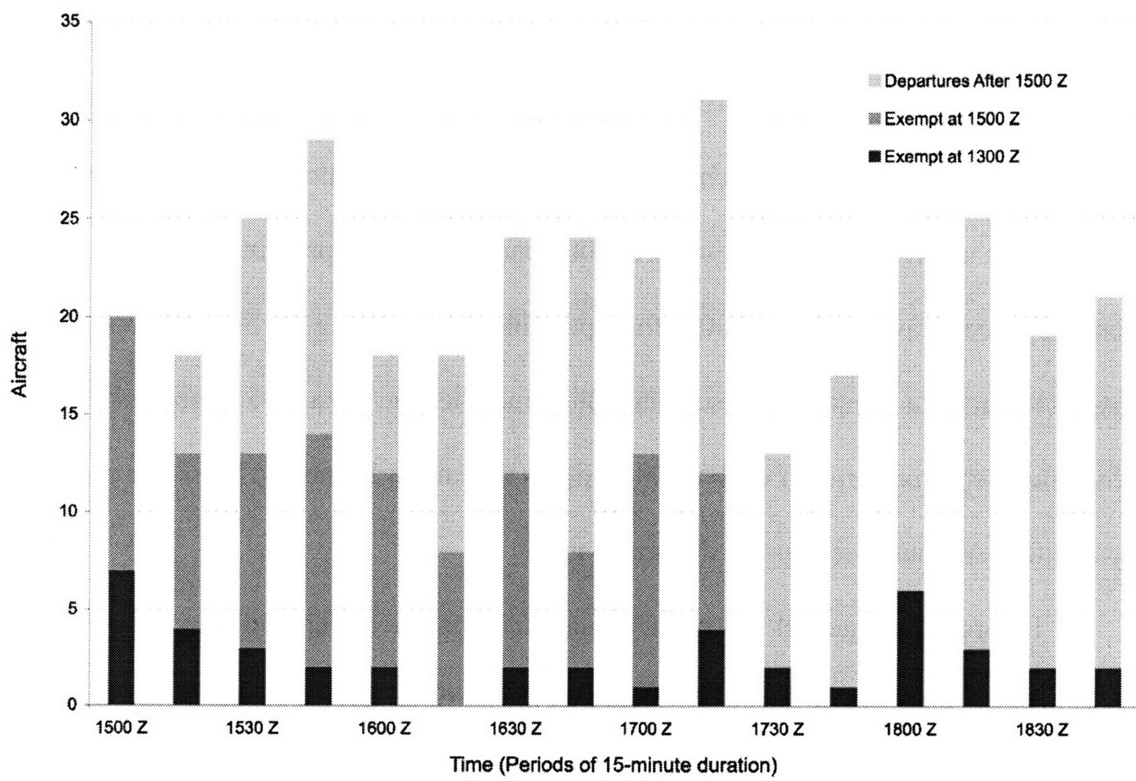

[12] The FAA has restricted the number of scheduled flights to 88 arrivals/hour since 2004

Flight Exemptions

The list of scheduled flight arrivals includes exempt flights, which are those that originate from an international airport, are already airborne, or are otherwise designated as exempt by the traffic manager. For example, of the 348 flights in the sample demand list, 43 (12.4%) either originate at an international airport or are already airborne at 1300 Z. If ground delays are not assigned until 1500 Z, then the number of international and airborne flights increases to 141 (40.5%). Exempt flights are not eligible to receive ground delay but, like non-exempt flights, they may receive airborne delay; thus, total delay cost of a GDP should include the costs incurred by exempt aircraft. For examples discussed in this thesis, it is assumed that the set of exempt aircraft (at 1300 Z) is known *a priori* and that the exemption status of a flight *f* is indicated by a parameter $XMT_f \in \{0,1\}$, where $XMT_f = 1$ for all exempt flights and 0 for non-exempt flights.

As formulated, the ECM assumes that all flights are eligible to receive both ground and airborne delay. Therefore, two important modifications are made in order to consider the airborne delay cost to exempt aircraft without assigning them ground delay. First, each exempt flight in the set of flights F is replaced with a proxy for use by the ECM. Each proxy flight has the same scheduled arrival time as the flight it replaces but a scheduled departure time equal to time period $t = 0$. For consistency, the value of $XMT_f = 1$ for all proxy flights.

Second, a set of additional constraints (3.16) are added to the formulation that require the departure status variable of each flight in time period $t = 0$ to be equal to the value of its exemption status parameter. As $XMT_f = 1$ for the exempt/proxy flights, these constraints ensure that all proxy flights depart as scheduled, in time period $t = 0$. All non-exempt flights, for which $XMT_f = 0$, remain eligible to receive ground delay.

$$d_{f0_q} = XMT_f \qquad \forall\ f \in F;\ q \in Q \tag{3.16}$$

The solution to the ECM with these proxy flights is consistent with the use of flight exemptions. First, neither the proxy nor originally exempt flights receive ground delay. Second, as the proxy flights have the same scheduled arrival times as the flights that they replace, the cumulative arrival demand at the airport is unchanged. Third, the proxy flights may receive airborne delay and the cost of this delay is considered by the ECM when assigning ground delay to non-exempt flights. And, as an additional benefit, replacing the airborne flights with proxies that depart at $t = 0$ means that the ECM does not need to consider any time periods prior to 1500 Z, thus reducing the required number of periods and the size of the model[13].

Maximum Allowable Arrival Queue

The maximum allowable arrival queue W_{tq}^{MAX} is the largest number of aircraft that the ECM permits in the arrival queue at the end of every time period (3.13). This limit may vary over time and for each profile and is assumed to be an input provided by the traffic manager. For the simple examples discussed here, W_{tq}^{MAX} is set equal to a hypothetical value of 30 aircraft for all time periods and profiles.

---

[13] §3.2.5

Functions of Delay Cost

Each experiment is based upon a set of common delay cost functions for ground, airborne, and cumulative delay. This common set serves two purposes. First, it is used as the basis for the objective function of the ECM, as well as those of other models from the literature to which the ECM is compared. Second, as the various models have different objective functions, this common set will also be used to evaluate the performance of each solution.

Figure 3-9 shows functions for the incremental delay cost of airborne, ground, and cumulative delay for each flight. The incremental cost of airborne delay increases with the duration of delay. For example, the cost of the first period of airborne delay is 1 unit and the fourth period is 11 units. Incremental ground delay cost is 0 units/period for the first eight periods of delay and 1 unit/period thereafter.

Figure 3-9 Functions for flight delay cost

t	α_t	γ_t	χ_t
0	0	0	0
1	1	0	1
2	3	0	1
3	5	0	1
4	11	0	1
5	11	0	1
6	11	0	1
7	11	0	1
8	11	0	1
9	11	1	1
10	11	1	1

Lastly, incremental cumulative delay cost has a value of 1 unit for all periods of delay time.

Recall that the delay cost model presented in §3.1 requires that the functions of ground, airborne, and cumulative delay cost be mutually exclusive in terms of the types of costs they represent. The set of functions in Figure 3-9 is consistent with this model. For example, the total delay cost for a flight that receives one period of ground delay and no airborne delay is 1 unit (Figure 3-10, page 64), which might represent the cost due to labor. The cost for a flight that receives one period of airborne delay and no ground delay is 2 units, representing labor and fuel. Lastly, the total delay cost for the flight in Figure 3-4 (page 48), which receives one period of ground delay and two periods of airborne delay, is 7 units: three periods of labor, two of fuel, and an additional penalty to account for the possibility of a diversion.

Figure 3-10 A Sample Calculation of the Total Delay Cost for a Flight

Flight	Ground Delay	Airborne Delay	Expression of Total Cost	Total Cost
f_1	1	0	$\gamma_1 + \chi_1$	= 1
f_2	0	1	$\alpha_1 + \chi_1$	= 2
f_3	1	2	$\gamma_1 + \alpha_1 + \alpha_2 + \chi_1 + \chi_2 + \chi_3$	= 7

Although the cost functions in Figure 3-9 are hypothetical, the increasing incremental costs are suggestive of how delay costs may be incurred in practice. For example, the cost of a flight diversion, which becomes more likely as the duration of airborne delay increases, is represented by an increase in the incremental cost of airborne delay. Furthermore, the increase in the incremental ground delay cost might represent the cost of a flight cancellation that may occur due to ground delays of extended duration.

The Distribution of Delay Among Aircraft

It is possible to affect how the ECM distributes delay among flights by modifying the cost functions used by the model. Specifically, if the marginal cost of delay increases with the duration of delay, then the ECM will favor solutions that distribute delay more evenly over those that concentrate the assignment of delay to a few flights. However, this is not true in the functions shown in Figure 3-9 for all t. Therefore, the cumulative delay cost function is adjusted for use in the ECM to include a super-linear function. The cost functions shown in Figure 3-11 are identical to those in Figure 3-9, except the parameters for incremental cumulative delay cost are replaced with:

$$\chi_t = t^{1+\varepsilon} - (t-1)^{1+\varepsilon}, \text{ where } \varepsilon = 1.0\text{E-5}$$

Figure 3-11 Delay Cost Parameters Used in the ECM

t	α_t	γ_t	χ_t
0	0	0	0
1	1	0	1
2	3	0	1 + 13E-6
3	5	0	1 + 19E-6
4	11	0	1 + 22E-6
5	11	0	1 + 25E-6
6	11	0	1 + 27E-6
7	11	0	1 + 28E-6
8	11	0	1 + 30E-6
9	11	0	1 + 31E-6
10	11	0	1 + 32E-6

The slight increase in χ_t with t in Figure 3-11 causes the formulation to favor solutions that distribute delay more evenly. Consider a simple example with two flights f_1 and f_2 and two solutions. In solution #1, f_1 receives one period of cumulative delay and f_2 three periods, and in solution #2, both flights receive two periods of delay. Although both solutions have an equal number of cumulative delay periods, solution #2, which assigns delay more evenly, will have a lower value as given by the super-linear function (Figure 3-12).

Figure 3-12 Cumulative delay cost example[14]

Possible Solution	**f_1 delay** (periods)	**f_2 delay** (periods)	**Cost (Fig 3-09)**	**Super-linear Cost (Fig 3-11)**
#1	1	3	4	4 + 32 E-6
#2	2	2	4	4 + 26 E-6

[14] $(\chi_1) + (\chi_1 + \chi_2 + \chi_3) = 1.0 + 3.000032959$

Comparison to Prior Models in the Literature

Within the SAGHP literature, the two closest models to the ECM are the Richetta-Odoni partially-dynamic LP model (RO-PDM) and the Mukherjee fully-dynamic MIP model (M-DM). As discussed in §2.3.3, Mukherjee improves upon Richetta-Odoni by allowing for the early release of aircraft when the arrival capacity forecast is revised upward. Neither of these models are able to consider flight-specific, non-linear airborne delay cost functions, such as those in Figure 3-9, which are assumed to represent the actual cost of delays in practice. Instead, both require a parameter that represents the ratio of the cost of one period of airborne delay to that of one period of ground delay. For the examples that are discussed in this thesis, both M-DM and RO-PDM are solved subject to a linear delay cost function, where ground delay has a cost of 1 unit/period and airborne delay of 2 units/period.

Mukherjee also proposes alternative objective functions for M-DM that yield solutions that distribute delay more evenly across aircraft. One alternative (M-DMDD) amends the objective function to include a cost penalty that increases with the square of ground delay time. Therefore, for the examples discussed in this thesis, the solution provided by the ECM is compared to those from three other models:

1. Mukherjee Dynamic Model (M-DM)
2. Mukherjee Dynamic Model with Distributed Delays (M-DMDD)
3. Richetta-Odoni Partially-Dynamic Model (RO-PDM)

Figure 3-13 briefly summarizes the models and their qualities.

Figure 3-13 A comparison of various SAGHP models

Model	Dynamic Assignment	Non-Linear Costs	Distributes Delays	Applies Cost Function
ECM	Fully	Cumulative, Ground, Air	Yes (Cumulative)	Super-linear χ_t
M-DM	Fully	None	No	Air/Ground (A/G) = 2.0
M-DMDD	Fully	Ground	Yes (Ground)	A/G = 2.0; plus $(\text{Delay})^2$
RO-PDM	Partially	None	No	A/G = 2.0

As the various models minimize different functions of delay cost, a direct comparison of the respective objective function values is meaningless. Therefore, as a basis for comparing the models, the cost functions proposed in Figure 3-9 are applied to each solution. Throughout the discussion in this thesis, the "cost" of a solution refers to the value of the under the common cost functions, while "objective function value" refers to the value of the objective function as specified in Figure 3-13.

Implementation

All models were implemented in OPL Studio and solved using CPLEX (v. 09) on a Pentium 4 desktop computer. The ECM and both Mukherjee models were solved as MIPs, and the Richetta-Odoni formulation as an LP.

Section 3.3.2 A Deterministic Example

The first example considers a deterministic arrival capacity forecast, which is represented as a scenario with a single, nominal capacity profile. This deterministic example demonstrates that the ECM yields a reasonable result for a nominal airport capacity rate and also offers a convenient starting point for comparisons between the ECM and other models from the literature. Furthermore, this deterministic example serves as a control for additional experiments with exhibit stochastic and dynamic arrival capacity.

Arrival Capacity Scenario

The arrival capacity scenario in the deterministic example consists of a single profile with a capacity of 25 arrivals in each 15-minute period, a rate equivalent to 100 arrivals/hour. This arrival capacity rate is on a par with the nominal operating capacity of a few large, metropolitan airports, such as Chicago's O'Hare International Airport (ORD). The average demand over the four-hour time horizon is well within the capacity of the airport. However, spikes in demand exceed capacity at 1545 and 1715 Z (Figure 3-14), which suggests that there will be some airborne arrival delays if no GDP is implemented.

Figure 3-14 Arrival capacity scenario for a deterministic example

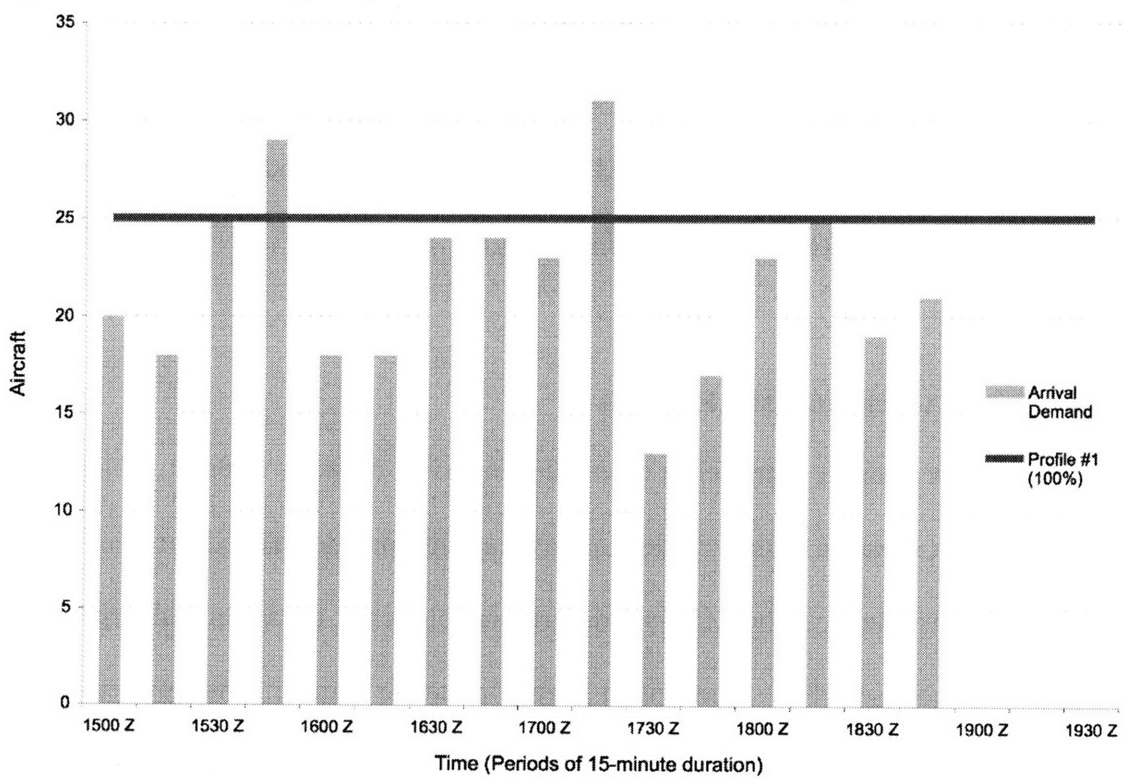

It is assumed that the time at which a GDP might start is 1500 Z and that the current time, at which the GDP might be initiated, is 1300 Z. The initiation time is relevant because all flights that are scheduled to depart before 1300 Z (and thus already airborne when the decision is made) are exempt from ground delay, regardless of their scheduled arrival time.

Results

The SAGHP is solved for the deterministic example using five models: the four noted previously plus a fifth model that gives the delay cost if no GDP is implemented. The fifth model (No-GDP) is identical to the ECM, except that no ground delays are permitted. The results of the models are presented in Figure 3-15 (page 70), which lists, for each solution, the:

1. Total cost (in units)
2. Total ground and airborne delay (in periods)
3. Length of the maximum observed airborne queue (in aircraft)
4. Maximum ground, airborne, and cumulative delay (in periods) observed to occur for any one flight
5. Time required to run each model (in seconds), as well as the number of variables and constraints in the CPLEX implementation

The ECM required 1.5 seconds to find an optimal solution for the deterministic example. This solution assigns one period of ground delay to each of ten aircraft, for a total of ten periods of ground delay; no airborne delay is observed. As compared to the

ECM, No-GDP results in the same amount of total delay time, except that all delay occurs in the air instead of on the ground. Thus, for this example, implementing the GDP suggested by the solution of the ECM reduces the total delay cost by 50%, as the air delay that would occur without a GDP is replaced by an equal amount of ground delay.

The other models, M-DM, M-DMDD, and RO-PDM, result in optimal solutions with the same cost, ground delay, airborne delay, and observed maximum flight delays and arrival queues as the ECM. These solutions are said to be *equivalent*, which is defined as meeting the following four criteria:

1. Assigning the same total amount of air and ground delay
2. Having the same maximum observed arrival queue length
3. Resulting in the same maximum observed flight delay
4. Have the same total expected cost

Figure 3-15 Results of the first, deterministic example

	ECM	M-DM	M-DMDD	RO-PDM	No-GDP
Total					
Cost	10	10	10	10	20
Ground Delay	10	10	10	10	0
Air Delay	0	0	0	0	10
Max. Queue	0	0	0	0	6
Flight Max					
Ground Delay	1	1	1	1	0
Air Delay	0	0	0	0	1
Cum. Delay	1	1	1	1	1
Model Stats					
Runtime (sec)	1.5	0.79	0.88	0.09	1.24
Variables	46,007	28,229	28,577	5,829	
Constraints	76,084	37,809	38,157	3,072	

Units	Cost:	units of cost	Delay:	periods
	Max. Queue:	aircraft		

Note that equivalent solutions may differ in the specific assignment of ground delay to flights, which means that there may exist multiple optimal and equivalent solutions for a given problem. For the purposes of evaluating the various models, these flight-specific differences are not considered.

That the solution to the ECM is equivalent to those of prior models in the literature is expected. The key improvements associated with the ECM become evident when there is a possibility of lengthy or numerous airborne delays, neither of which are present in this example.

Section 3.3.3 Low Likelihood of a Severe Capacity Reduction

The second example is motivated by the design of GDPs in practice, where the objectives may include avoiding outcomes with excessive airborne delay. This example features a hypothetical capacity scenario with a low but appreciable likelihood of a severe reduction in capacity. If the severe reduction is realized, it might lead to substantial delays, flight diversions, and an unacceptable number of aircraft in the airborne arrival queue if a GDP is not implemented proactively. However, the probability of the severe event is low, so that models that seek only to minimize the total *expected* delay cost may not assign sufficient ground delay to avoid an unacceptable outcome.

Arrival Capacity Scenario

The arrival capacity scenario in this example has two profiles. The first profile represents the nominal airport arrival capacity (25 aircraft/period), which is identical to that of the previous example in §3.3.2. For the second profile, which has a likelihood of

10%, the airport arrival capacity rate is temporarily reduced to 0 aircraft/period at 1500 Z, maintains this rate for four periods (one hour), and then returns to the nominal rate. This hypothetical example might represent the passing of a weather front, where there is a small likelihood that airport operations will cease due to severe weather[15]. For each profile, the capacity of the arrival queue is 30 aircraft.

For this second example, the GDP is modeled dynamically; it is assumed that an initial decision of whether or not to initiate a GDP is made at 1300 Z and that the assignment of delay may be revised once, at 1500 Z. Figure 3-16 shows the arrival capacity scenario superimposed on the cumulative scheduled arrival demand. The columns indicating cumulative arrival demand are separated into three parts (denoted by the dark, medium, and light shading), which represent the flights that are exempt at 1300 Z, the flights that would be exempt at 1500 Z if no proactive action is taken, and flights that are scheduled to depart after 1500 Z.

Results

In this second example, the solution to the ECM initially (at 1300 Z) assigns a total of 144 periods of ground delay to 94 aircraft. The flights delayed as part of this initial GDP represent 96% of the 98 non-exempt departures scheduled between 1300 Z and 1500 Z. Most of these aircraft are held for a short period of time and then released; by the time of the forecast revision at 1500 Z, only 16 of the 94 delayed aircraft remain on the ground. The average assigned delay per delayed aircraft is 1.53 periods, or about 23 minutes.

---

[15] The scenario tree for this example is the same as that shown previously in Figure 3-3, page 45

Figure 3-16 Arrival capacity and demand for the second example

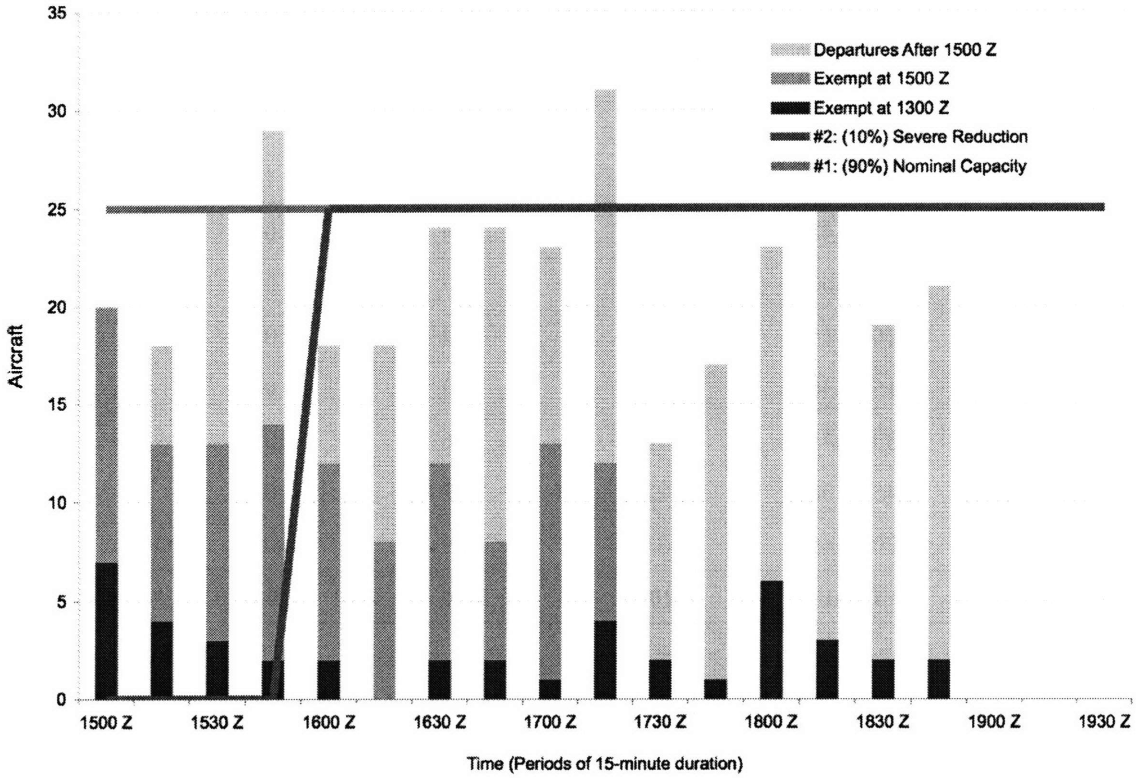

At 1500 Z, the assignment of ground delay is revised. If Profile #1 is realized, which indicates nominal capacity, all of the aircraft that are being delayed on the ground are immediately released. In this case, through the remainder of the program, an additional 29 aircraft would each receive one period of ground delay. Alternatively, if Profile #2 is realized, the 16 aircraft already being delayed on the ground are assigned additional delays of up to one hour. Furthermore, under Profile #2, a total of 745 periods of ground delay are also assigned to the 207 flights that are scheduled to depart after 1500 Z, for an average of 3.6 periods, or 54 minutes, of ground delay per flight (Figure 3-17).

Figure 3-17 Ground delay assigned under Profile #2 for departures after 1500 Z

Departure Time	Count of Flights	Avg. Ground Delay/Flight	
		Periods	Time (h:mm)
1500-1515 Z	42	4.85	1:13
1515-1530 Z	21	4.38	1:06
1530-1545 Z	17	3.94	0:59
1545-1600 Z	22	3.77	0:57
1600-1615 Z	16	3.44	0:52
1615-1630 Z	18	3.22	0:48
1630-1645 Z	18	3.06	0:46
1645-1700 Z	22	2.72	0:41
1700-1845 Z	31	2.29	0:34
Total	**207**	**3.60**	**0:54**

Despite the significant ground delays assigned by the ECM, the solution still does not avoid substantial airborne delays. If Profile #2 is realized, 51 aircraft receive a total of 116 periods of airborne delay (2.27 periods/delayed aircraft), with a total airborne delay cost of 358 units (7.02 units/delayed aircraft). The maximum observed size of the airborne queue under Profile #2 is 30 aircraft (Figure 3-18), which indicates that the constraints (3.14) are binding in the optimal solution.

For this example, the cost of the solution to the ECM is greater than that of the other models (Figure 3-20, page 77). The M-DMDD solution has the lowest cost (202.4 units), while the M-DM, and RO-PDM solutions cost slightly more. The cost of the ECM solution is 296.5 units, a 46.5% increase over M-DMDD. However, in comparison to the other models, the ECM is the only one to avoid an unacceptably large airborne queue under Profile #2. The other solutions could result in as many as 56 aircraft in the airborne queue, nearly twice the number as the ECM.

Figure 3-18 The airborne queue over time for the solution to the ECM

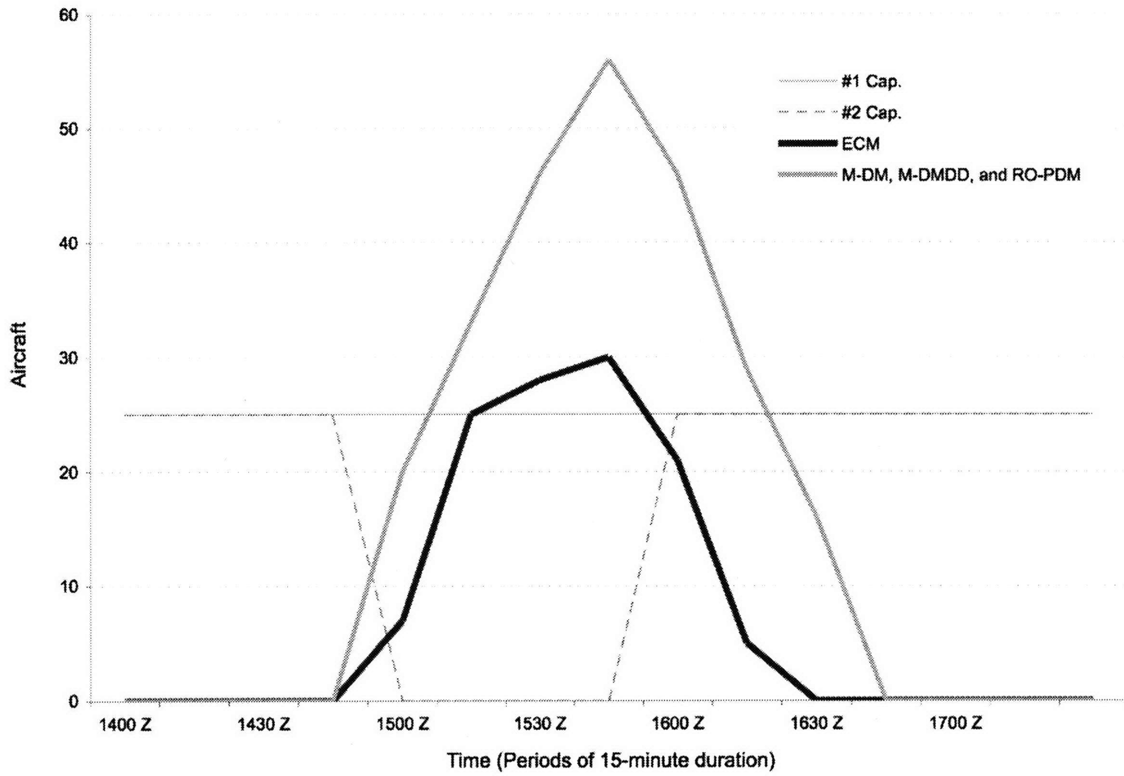

The total expected cost of the ECM is greater because the maximum queue constraints (3.15) require that additional flights be proactively delayed on the ground. For example, in contrast to the 94 flights that are assigned ground delay prior to 1500 Z by the ECM, only 10 are delayed proactively by M-DMDD. Furthermore, in order to accommodate the arrival of the flights delayed at 1300 Z, additional ground delay (and cost) is incurred after 1500 Z. If Profile #1 is realized, the solution to the ECM results in an additional 29 periods of ground delay after 1500 Z while the Mukherjee solution requires none! Thus, while the ECM reduces both the cost and observed airborne queue of the worst-case outcome, it increases the cost of the best, most likely case.

The ECM is used to solve the problem under three alternative hypotheses. The first (No-MQC) solves for an assignment of ground delay for which there are no constraints on the size of the arrival queue. The expected cost of the No-MQC solution is

32.4% lower than that for the ECM and 1% lower than M-DMDD (Figure 3-21, page 78). The No-MQC result shows that in the absence of the airborne queue size constraint the ECM identifies a lower cost solution than Mukherjee – the key difference between No-MQC and M-DMDD being the direct consideration of non-linear cost functions.

The second and third provide an upper and lower bound on delay costs. The upper bound is a solution in which no ground delays are assigned (No-GDP) and that results in a total expected cost of 489.6 units. The lower bound is found by assuming that the initial assignment of ground delay is based upon "perfect information" (PI), meaning that there is no uncertainty in the 1300 Z arrival capacity forecast[16]. For this example, perfect information results in a total expected cost of 133 units. Figure 3-19 summarizes the solutions of the different models.

Figure 3-19 A comparison of the costs for the solution to each model

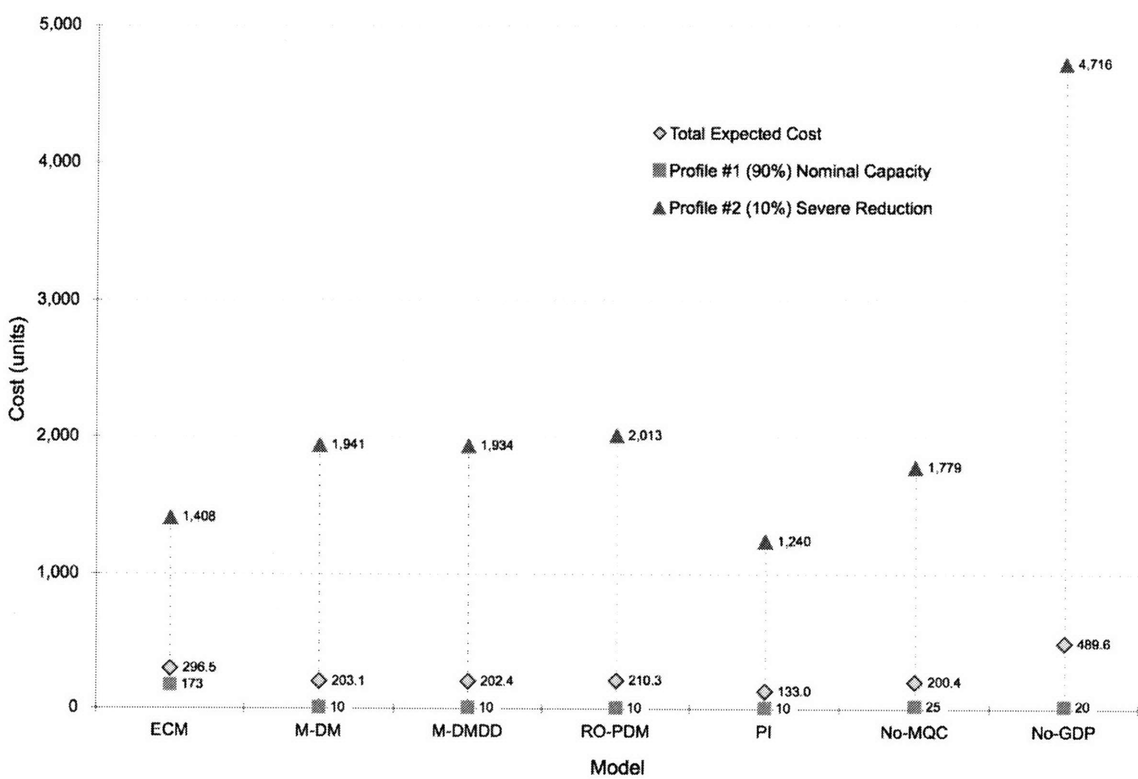

[16] The PI model is identical to the ECM except the coupling constraints (3.15) are removed.

	ECM			M-DM			M-DMDD			RO-PDM		
Total	Exp	#1	#2	Exp	#1	#2	Exp	#1	#2	Exp	#1	#2
Cost	296.5	173	1408	203.1	10	1941	202.4	10	1934	210.3	10	2013
Ground Delay	249	173	934	89	10	804	89	10	804	89	10	801
Air Delay	12	0	116	25	0	246	25	0	246	25	0	249
Max. Queue	30	0	30	56	0	56	56	0	56	56	0	56
Flight Max												
Ground Delay	6	5	6	10	1	10	7	1	7	15	1	15
Air Delay	4	0	4	4	0	4	4	0	4	4	0	4
Cum. Delay	6	5	6	10	1	10	7	1	7	15	1	15

Model Stats	ECM	M-DM	M-DMDD	RO-PDM
Runtime	7.05	1.77	1.92	0.22
Variables	110,185	56,457	57,153	11,658
Constraints	188,867	78,749	79,445	9,988

Figure 3-20 A comparison of the solution of the ECM to other models

	ECM			No GDP			PI			No-MQC		
Total	Exp	#1	#2	Exp	#1	#2	Exp	#1	#2	Exp	#1	#2
Cost	296.5	173	1408	489.6	20	4716	133	10	1240	200.4	25	1779
Ground Delay	249	173	934	0	0	0	109	10	1002	105	25	825
Air Delay	12	0	116	114	10	1050	5	0	48	23	0	225
Max. Queue	30	0	30	92	6	92	16	0	16	54	0	54
Flight Max												
Ground Delay	6	5	6	n/a	n/a	n/a	5	1	5	7	4	7
Air Delay	4	0	4	4	1	4	4	0	4	4	0	4
Cum. Delay	6	5	6	4	1	4	5	1	5	7	4	7
Model Stats												
Runtime	7.05			3.33			5.09			5.73		
Variables	110,185											
Constraints	188,867											

Figure 3-21 Additional results for the second example

As a final note, in regards to runtime, the ECM also takes longer than other models. However, at 7.05 seconds, the run time of the ECM would be acceptable for use in practice. Run time of the ECM will be discussed in greater detail in §4.3.2.

Section 3.3.4 Conclusions from the Two Examples

The two examples in §3.3.2 and §3.3.3 illustrate some of the differences between the various models. In the first example, the ECM finds a solution equivalent to that of the Mukherjee and Richetta-Odoni models; in the second, the ECM solution is more costly, but also more practical. Note that the two examples are similar: both use the same flight demand list and assume the same decision times. The only difference between the two examples is that, in the first, the nominal capacity profile is guaranteed, while in the second, there is a small likelihood of a severe storm. The ECM is capable of addressing both examples, while the solutions to the Mukherjee and Richetta-Odoni models, which lack constraints on the length of the arrival queue, might result in an unacceptable outcome in the second example.

The ECM avoids extensive airborne delays because it assigns ground delay proactively. The size of the departure queue, or the number of aircraft in a ground hold, is never more than two aircraft in the first example. However, under Profile #1 for the second example, the departure queue is as large as 30 aircraft (Figure 3-22, page 80). On the other hand, M-DMDD assigns nearly the same amount of ground delay under Profile #1 in the second example as it does in the first example (Figure 3-23, page 81) – in the second example, the M-DMDD effectively takes no proactive action! The Mukherjee model minimizes the expected cost and, as a result, does not assign the additional, preemptive delays that are needed to reduce the observed airborne queue for Profile #2 to an acceptable level.

In contrast to prior models in the SAGHP literature, the ECM is better able to identify the risk posed by the possibility of a severe storm and takes proactive action to assign delays. Furthermore, a feature of the ECM is that the degree to which the model is averse to risk can be adjusted. As the results of the No-MQC test indicate, the queue size constraints (3.13) can be relaxed so as to not be binding in the optimal solution. In this case, the ECM finds a solution (No-MQC) with a cost that is lower than the other SAGHP models.

Figure 3-22 Departure queue over time (ECM)

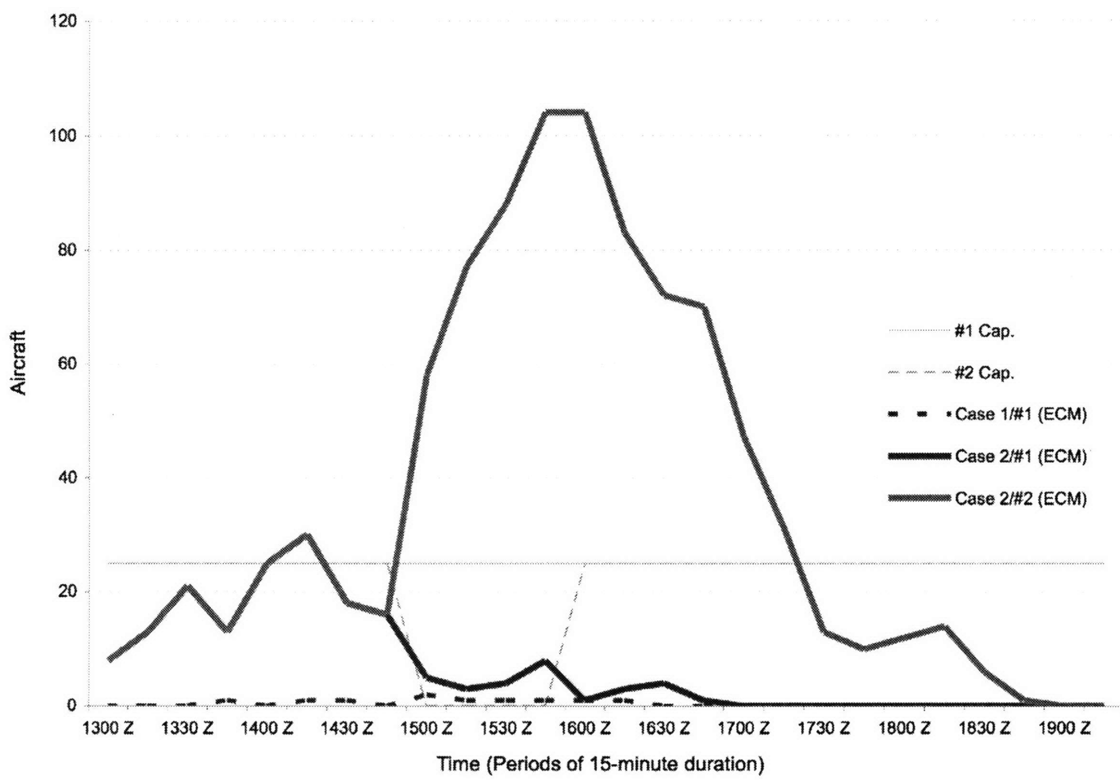

Figure 3-23 Departure queue over time (M-DMDD)[17]

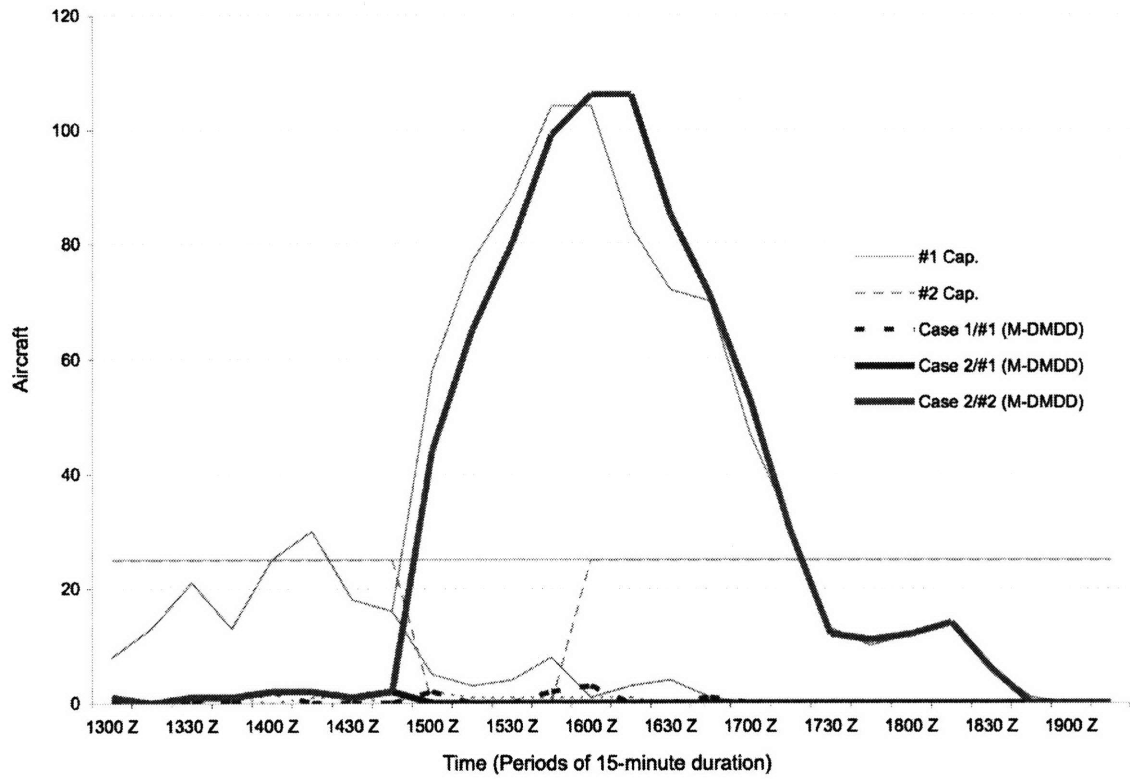

As a final observation, note that the linear cost functions used in the Mukherjee formulations could also be easily applied to the ECM. With identical cost functions and no restriction on the airborne queue, the ECM would yield a solution equivalent to that of Mukherjee in the second example, as it did in the first example. Thus, the ECM is not only more applicable, but it is also a more general formulation than Mukherjee, as well.

Section 3.4 Conclusions

A GDP is a tool used by traffic managers at the FAA to manage the arrival demand at capacity-constrained airports. GDPs reduce the amount of airborne congestion by proactively delaying flights before they depart. However, a program must often be

[17] For comparison, the departure queue for the solution to the ECM is included, as well, and marked in thin gray lines.

implemented before the future arrival capacity of the capacity-constrained airport is known with certainty. The decision of which flights to delay, and by how much, is described in the academic literature as the single-airport ground hold problem, or SAGHP.

Chapter Three proposes a new model to address the SAGHP: the ECM, or "extended cost model." As compared to previous models, the ECM generalizes the calculation of delay costs, permitting the incremental cost of a delay to vary with its duration. The model also considers an additional GDP objective, which is to prevent the formation of lengthy airborne arrival queues that might be undesirable or unacceptable to traffic managers. In comparison to other models, the ECM is able to incorporate more of the factors that are considered in practice.

The ECM compares favorably, in many respects, to other models when applied to two hypothetical examples. In the second example, for which the arrival capacity is both uncertain and dynamic, and for which there is a low, but appreciable likelihood of a severe reduction in arrival capacity, the ECM is solved with and without restrictions on the allowable size of the arrival queue. With the restriction, the solution assigns delays more aggressively than other models and reduces the possible size of the arrival queue significantly. Without the restriction, the solution has a lower cost than any of the other models tested.

In addition to demonstrating some of the advantages of the ECM over other models in the literature, Chapter Three also suggests how the model might be applied in practice. Assuming the availability of an arrival capacity scenario, the ECM is adapted to use ETMS arrival demand data, as well as account for flights exempt from ground delays.

Practical considerations impose additional requirements on model performance: the solutions should be robust in regards to uncertainty in the input parameters (such as cost, capacity, etc.) and the formulation should solve sufficiently quickly so as to be of use during the design of a GDP in practice. Chapters Four and Five explore the performance of the ECM under a wide variety of examples in order to better demonstrate its capabilities.

Chapter 4 An Analysis of the Sensitivity of the Solution to the ECM to Changes in the Values of Input Parameters

The ECM, or "Extended Cost Model," solves the single airport ground hold problem (SAGHP) subject to uncertain and dynamic arrival capacity. As with other SAGHP models, the ECM is motivated by the FAA's use of Ground Delay Programs to manage the arrival demand at capacity-constrained airports. However, unlike previous models, the ECM also considers non-linear airborne delay cost functions and the capacity of the airborne arrival queue, thus avoiding solutions that might be unacceptable in practice. For two examples discussed in Chapter Three, solutions to the ECM compare favorably to those from previous models in the literature.

An example in the previous chapter applies the ECM to a problem in which there is a low, but appreciable likelihood of a severe reduction in arrival capacity. Such an example raises many interesting questions about how the assignment of delay would change under various alternative assumptions. For example, what if the reduction in arrival capacity is less severe? What if the reduction is more likely to occur? Or, what if the capacity of the airborne arrival queue is greater? Examining how the optimal solution changes under alternative assumptions also provides insight into how the ECM compares to alternative models in the literature, as well as how it might be applied in practice.

Chapter Four presents the results of a sensitivity analysis of the solution to the ECM to changes in the input parameters. First, §4.1 provides an overview of the objectives of the analysis and describes the experimental setup. Then, §4.2 discusses the results of nine separate experiments that test the sensitivity of the solution to changes in various input parameters. Finally, §4.3 draws conclusions about the performance of the ECM for solving the SAGHP and for use in practice.

Section 4.1 Objectives of the Sensitivity Analysis

The advantages of the ECM in comparison to models in the SAGHP literature may be viewed from two different perspectives. First, from the perspective of the literature, the ECM generalizes prior models to consider a wider range of delay cost functions and the capacity of the airborne arrival queue. Second, from the perspective of an air traffic manager, the ECM minimizes the cost of delays subject to constraints that prevent solutions that might permit an unacceptably large airborne queue.

Correspondingly, the sensitivity analysis presented in this chapter seeks to provide insight for both of these perspectives. The analysis has three primary objectives, which are to: 1) demonstrate that the ECM can consider various input parameters with a wide range of values; 2) show that the ECM solves non-trivial problems quickly enough to make it useful in practice; and 3) identify the relationship between various input parameters and the optimal assignment of ground delay.

Section 4.1.1 Conducting the Analysis

The sensitivity analysis is conducted as nine separate experiments. Each experiment measures the change in value of various dependent variables (such as the total expected cost) in response to a change in one or more independent variables. The independent variables are various input parameters to the ECM and are listed for each experiment in Figure 4-1.

Figure 4-1 A summary of experiments in the sensitivity analysis

Exp.	Independent Variable(s)	Range of Values	Count of Trials
1	Airport Arrival Capacity	0 – 30 aircraft	14
2	Profile Likelihood	0.00 – 1.00	21
3	Airborne Queue Capacity	15 – 60 aircraft	18
4	Time of Initial Decision	1300 – 1500 Z	9
5	Airborne Queue Capacity Time of Initial Decision	15 – 60 aircraft 1300 – 1500 Z	69
6	Time of Revision	1300 – 1600 Z	18
7	Airborne Delay Cost (A)	1 – 55 units	15
8	Airborne Delay Cost (B)	0 – 10	20
9	Airborne Delay Cost (C)	0 – 25 units	16

The sensitivity analysis highlights the conceptualization of stochastic and dynamic arrival capacity forecasts and the input parameters that are unique to the ECM among models in the SAGHP literature. The experiments focus on three areas:

1. The arrival capacity scenario tree (experiments #1, 2, 4, 5, and 6)
2. The capacity of the airborne arrival queue (3 and 5)
3. The cost of airborne delay (7, 8, and 9)

Section 4.1.2 Experimental Setup: The Base Case

The nine experiments share a common experimental setup that is previously presented as an example in §3.3.3. For each experiment, arrival demand is given by a list of 348 flights that are scheduled to land at Chicago O'Hare International Airport over a four-hour period of time. The arrival capacity scenario tree contains two profiles, one that represents the nominal arrival rate (25 arrivals/period) and the other a severe reduction (0 arrivals/period). A comparison of the arrival demand and capacity under each profile for the base experimental setup is shown in Figure 3-16 (page 73). The scenario tree contains two nodes, or decision times. At 1300 Z, a GDP is initiated for a forecast specifying that the likelihood of Profile #1 is 0.90 and that of Profile #2 is 0.10. At 1500 Z, the actual arrival capacity of the airport is revealed and the GDP is revised accordingly.

For these experiments, it is assumed that the cost of delay to each flight is represented by the functions in Figure 3-9. However, as discussed in §3.3.1, these functions are modified slightly for use by the ECM in order to identify solutions that distribute delay more evenly among aircraft. Please note that the term "cost" as used in this chapter refers to the value obtained by applying the modified functions in Figure 3-11 to the solution of the ECM.

This example, which is summarized in Figure 4-2, is used as the basis for the sensitivity analysis for two reasons. First, with two contrasting profiles, the example is simple enough to allow meaningful conclusions to be drawn about the sensitivity of the solution to the individual input parameters, yet complex enough to represent the core tradeoff in the design of a GDP between the control of flights and the availability of information. Second, this example is one for which the solution to the ECM holds some advantages over those of other models in the literature[18].

---

[18] §3.3.3

Figure 4-2 A summary of the base experimental setup

	Initial Capacity	Likelihood
Profile #1	25	0.90
Profile #2	0	0.10

	Initiation	Revision
Decision Time	1300 Z	1500 Z

	Airborne	Ground	Cumulative
Incremental Delay Cost	α_t	γ_t	χ_t

Arrival Queue Capacity	30 aircraft

Section 4.1.3 Terminology

To facilitate a discussion of the results of each experiment, several terms are defined below. Each term is defined with reference to the optimal solution of the ECM for a single trial.

- Cost under a profile: the sum of the delay costs across all flights assuming the realization of a particular arrival capacity profile
- Total expected cost: the expected value of sum of delay costs across all flights
- Marginal cost or marginal delay: the observed incremental change in cost or delay per unit of change in the independent variable with respect to the previous trial
- Maximum observed arrival queue: the maximum number of flights that would simultaneously experience airborne delay
- Maximum observed flight delay: the greatest amount of airborne, ground, or cumulative delay assigned to any one flight

Section 4.2 Results of the Sensitivity Analysis

In this section, the sensitivity analysis is presented as nine separate experiments. The general results of the sensitivity analysis as a whole will be discussed in §4.3.1. Detailed results are also provided in the appendix.

Section 4.2.1 Severity of the Capacity Reduction: M_{2t}

The first experiment varies the planned airport arrival capacity rate (PAAR) of Profile #2 between 1500 and 1600 Z, which is referred to as M_{2t}. The values of M_{2t} vary between 0, which is the value in the base setup, and 30 aircraft, which exceeds the nominal arrival capacity. A superposition of the profiles used in this experiment is shown in Figure 4-3. The arrival capacity under Profile #2 after 1600 Z and under Profile #1 for all time periods is held constant at 25 aircraft/period.

Figure 4-3 Arrival capacity profiles used in experiment #1

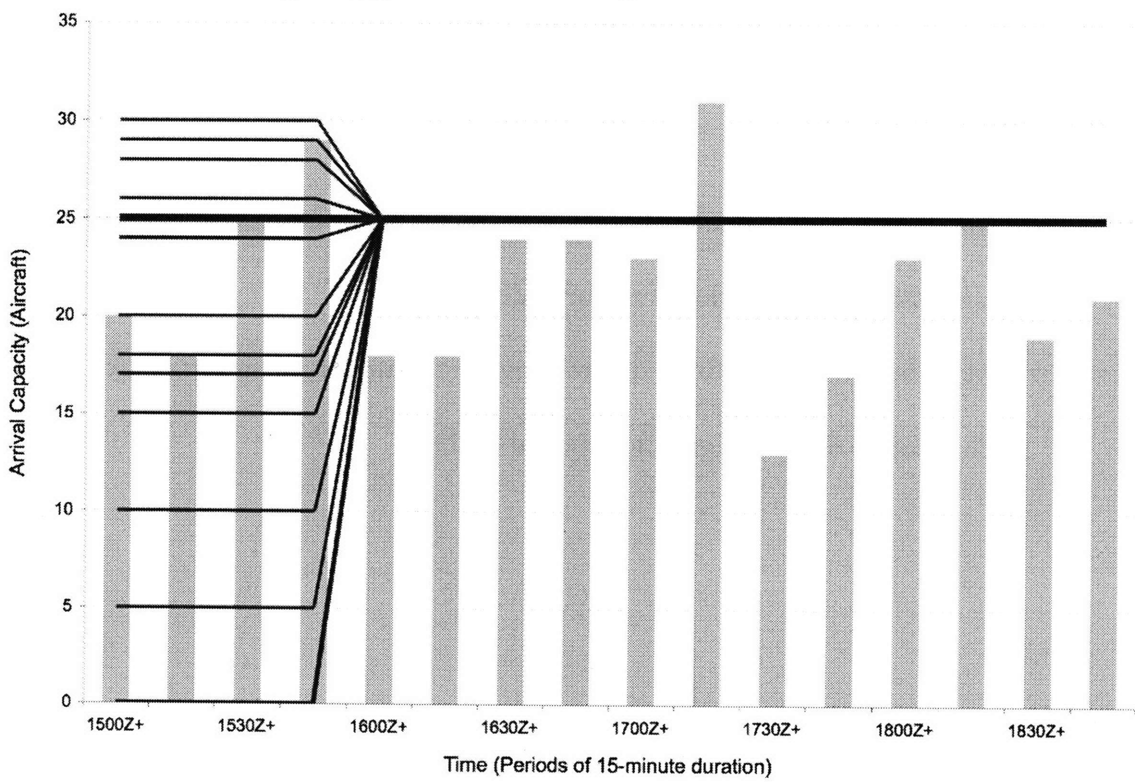

A summary of the setup for experiment #1 is shown in Figure 4-4. Note that the profile likelihoods are $\{p_1 = 0.10, p_2 = 0.90\}$. This change increases the sensitivity of the overall solution to the capacity of Profile #2.

Figure 4-4 A summary of experimental setup #1

	Initial Capacity	Likelihood
Profile #1	25	**0.10**
Profile #2	***Varies***	**0.90**

	Airborne	Ground	Cumulative
Incremental Delay Cost	α_t	γ_t	χ_t

Decision Time	Initiation	Revision
	1300 Z	1500 Z

Arrival Queue Capacity	30 aircraft

Results

In this experiment, as the arrival capacity under Profile #2 increases, the total expected cost decreases (Figure 4-5, page 92). However, the marginal benefit of an increase in capacity (measured as the reduction in total delay cost per unit change in capacity) is less for larger values of M_{2t}. For example, if M_{2t} is 0 aircraft, then increasing the capacity by 5 aircraft reduces the total expected cost by 459.5 units. However, if M_{2t} is 20 aircraft, then the same net increase in capacity reduces the cost by only 19.8 units. This suggests that an incremental increase in the arrival capacity or decrease in the arrival demand is most beneficial when the capacity is low or when the difference between demand and capacity is greatest.

Similarly, the cost under each profile also decreases as the capacity of Profile #2 increases. For example, the expected cost under Profile #1 decreases from 349.0 to 10.0 units as M_{2t} increases from 0 to 30 aircraft/period. The decrease in cost under Profile #1 is especially notable because the capacity of this profile is the same for all trials in the experiment; the reduction is due to the coupling constraints (3.15) that apply to all ground delay assigned prior to 1500 Z revision. Additional detailed results are available in the appendix.

Conclusions

This experiment examines the solutions of the ECM for 14 different values of the airport arrival capacity rate under Profile #2 between 1500 and 1600 Z. The results indicate that increasing the capacity under one profile reduces the total expected cost, as well as the cost under each profile. Furthermore, the marginal benefit of an increase if greatest when the arrival capacity is most constraining.

Figure 4-5 Delay cost as a function of M_{2t}

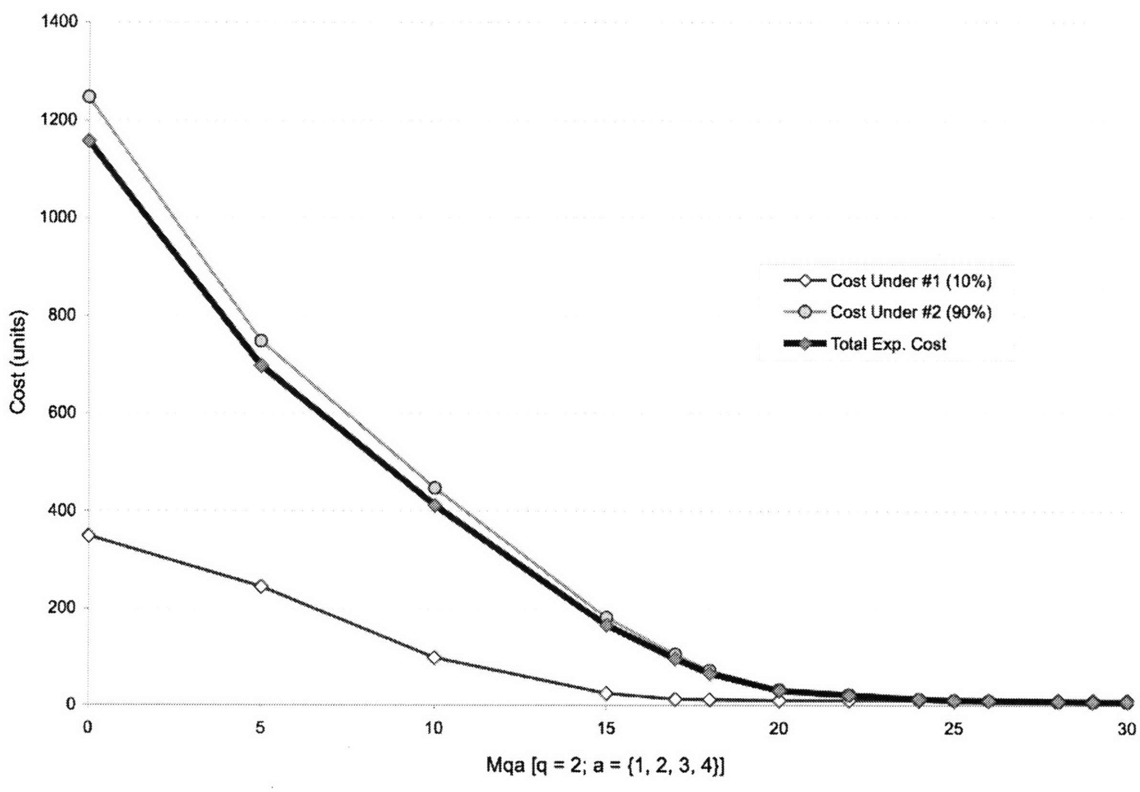

Section 4.2.2 Likelihood of a Reduction in Capacity: p_2

The second experiment varies the likelihood of the arrival capacity profiles (Figure 4-6). Recall that the cumulative likelihood $\sum_{Q} p_q = 1$; so for a scenario with two profiles, the likelihood of one profile can be expressed as a function of the other. For convenience, in this experiment, p_1 is treated as a dependent variable, with $p_1 = 1 - p_2$. Possible values of p_2 are specified on the interval $(0,1)$ in increments of 0.05.

Figure 4-6 A summary of experimental setup #2

	Initial Capacity	Likelihood
Profile #1	25	***Varies***
Profile #2	0	***Varies***

Decision Time	Initiation	Revision
	1300 Z	1500 Z

	Airborne	Ground	Cumulative
Incremental Delay Cost	α_t	γ_t	χ_t

Arrival Queue Capacity	30 aircraft

Of note in this experiment is that the scenarios in which $p_2 = \{0,1\}$ are equivalent to a scenario that contains only a single profile. For example, if $p_1 = 1$, then Profile #1 is assured of being realized; this represents the same forecast as the deterministic scenario used in the example in §3.3.2. As additional profiles with a likelihood of 0 do not change the forecast represented by a scenario, they are considered to be *irrelevant*.

Irrelevant profiles may affect the optimal solution to the ECM in two ways. First, the maximum queue constraints (3.13) apply to all profiles, without regard to the profile likelihood. Thus, an irrelevant profile may force the solution to consider the queue size of a profile that will not occur! This represents a limitation of the ECM. Second, if the likelihood of a profile is 0, then the cost of delays under that profile will not be factored in to the value of the objective function. As a result, the ECM may assign delays under irrelevant profiles that would be impractical or inconsistent with those for a profile with a positive likelihood. To avoid such solutions, the trials in this experiment solve the SAGHP at the limit as p_2 approaches 0 and 1. This will not affect the cost of the solution, but may change the delays assigned under an irrelevant profile.

Results

Figure 4-7 summarizes the results of the second experiment. As the likelihood of Profile #2 increases, the optimal solution assigns more ground delay, which increases the cost under Profile #1 and decreases the cost under Profile #2. Furthermore, as p_2 increases, the total expected cost curve is downward sloping (Figure 4-8)[19] [20]. This indicates that additional proactive ground delays reduce the cost under Profile #2 (weighted by p_2) by more than they increase the cost of Profile #1 (weighted by p_1).

The change in slope of the total expected cost curve is measured in terms of the change in the marginal cost, or second derivative. This is exhibited in Figure 4-7 as the column labeled "∂^2." Note that the trials for which the change in the marginal cost is equal to 0, such as those for $p_2 = [0.30, 0.50]$, correspond to solutions for which the assigned ground and airborne delays are identical to those of the next trial.

[19] The slope is highlighted by an additional straight line that is added to the chart.
[20] Please note that, for the charts exhibited in this chapter, a marker indicates the result of each trial and, where applicable, a straight line has been drawn between markers for successive trials to highlight the rate of change in value.

Figure 4-7 A summary of the results of experiment #2

			Profile #1				Profile #2				
p_2	Exp. Cost	∂^2	Cost	Air Delay	Ground Delay	Max. Queue	Cost	Air Delay	Ground Delay	Max. Queue	Run Time
0.00	160.0		160	0	160	0	1,581	151	899	30	9.26
0.05	229.8	-62	160	0	160	0	1,556	134	916	30	6.92
0.10	296.5	-99	173	0	173	0	1,408	116	934	30	7.15
0.15	358.3	-64	173	0	173	0	1,408	116	934	30	7.35
0.20	416.8	-42	189	0	189	0	1,328	100	950	30	7.40
0.25	473.3	-16	191	0	191	0	1,320	98	952	30	7.39
0.30	528.9	0	195	0	195	0	1,308	94	956	30	7.70
0.35	584.6	0	195	0	195	0	1,308	94	956	30	7.72
0.40	640.2	0	195	0	195	0	1,308	94	956	30	7.65
0.45	695.9	0	195	0	195	0	1,308	94	956	30	7.90
0.50	751.5	-14	195	0	195	0	1,308	94	956	30	8.12
0.55	806.5	-52	213	0	213	0	1,292	86	964	28	8.15
0.60	858.8	-24	233	0	233	0	1,276	78	972	26	7.98
0.65	910.0	0	245	0	245	0	1,268	74	976	25	7.97
0.70	961.1	-15	245	0	245	0	1,268	74	976	25	8.05
0.75	1,011.5	-2	257	0	257	0	1,263	71	979	24	8.20
0.80	1,061.8	-4	257	0	257	0	1,263	71	979	24	8.30
0.85	1,111.9	-96	301	0	301	0	1,255	63	987	20	8.09
0.90	1,157.2	-57	349	0	349	0	1,247	55	995	16	8.70
0.95	1,199.7	-42	433	0	433	0	1,240	48	1,002	16	9.80
1.00	1,240.0		433	0	433	0	1,240	48	1,002	16	10.18

Units	Max. Queue:	aircraft	Run Time:	seconds
	Exp. Cost, Cost:	units of cost	p_2:	likelihood of Profile #2
	Air, Ground Delay:	aircraft-periods	∂:	units of cost / change in p_2

Figure 4-8 Delay cost as a function of the likelihood of Profile #2

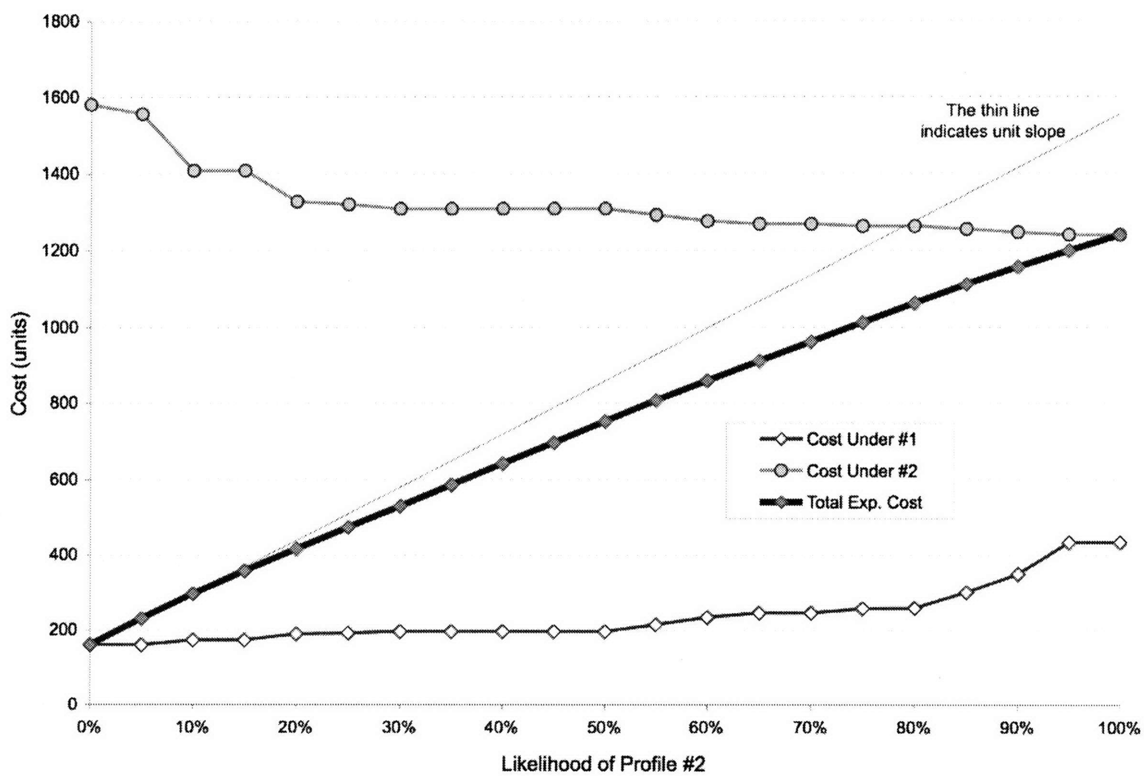

Conclusions

This experiment examines changes in the cost and assignment of ground delay for varying likelihoods of the arrival capacity profiles. Although an increase in p_2 generally decreases the cost under Profile #2 and increases the cost under Profile #1, there exist values of p_2 for which a net change of up to 0.2 does not change the assignment of ground delay in the optimal solution. For the design of GDPs in practice, this suggests that it may be useful to evaluate the sensitivity of the solution if the profile likelihoods, themselves, are uncertain. A more detailed examination of the sensitivity of the solution to changes in likelihood is an area of possible future research.

Furthermore, this experiment also demonstrates that solutions to the ECM may be sensitive to the addition of irrelevant profiles to the arrival capacity scenario. As noted previously, the scenario in which $p_1 = 0$ represents the same arrival capacity forecast as the scenario used in the example in §3.3.2. However, as the maximum queue (3.13) and coupling (3.15) constraints apply to the irrelevant profile, the solution to the ECM assigns more delay – and results in a higher total expected cost – for the scenario in this example than it does when there is only a single profile. Therefore, care should be exercised before using the ECM to ensure that the scenario tree and the values of W_{qt}^{MAX} are appropriate. Additional research might also explore alternative formulations in which the capacity constraints consider likelihood.

Section 4.2.3 The Capacity of the Airborne Arrival Queue: W_{qt}^{MAX}

The experiments discussed in the previous sections assume that the capacity of the airborne arrival queue W_{qt}^{MAX} is 30 aircraft. This experiment varies the arrival queue capacity between 16 and 60 aircraft. For values of less than 16 aircraft, the ECM does not have a feasible solution for this example. The experimental setup is summarized in Figure 4-9; note that the profile likelihoods are $\{p_1 = 0.95, p_2 = 0.05\}$.

Figure 4-9 A summary of experimental setup #3

	Initial Capacity	Likelihood
Profile #1	25	**0.95**
Profile #2	0	**0.05**

	Initiation	Revision
Decision Time	1300 Z	1500 Z

	Airborne	Ground	Cumulative
Incremental Delay Cost	α_t	γ_t	χ_t

Arrival Queue Capacity	***Varies***

Results

The maximum queue constraints are not binding for values of $W_{qt}^{MAX} > 56$. For these trials, the optimal solution assigns 10 periods of ground delay under Profile #1 and 804 periods of ground delay (plus an additional 247 periods of airborne delay) under Profile #2 (Figure 4-10, page 98). As the value of W_{qt}^{MAX} decreases, the constraints become binding and the optimal solution assigns more ground delay. Therefore, in this experiment, assigning additional ground delay decreases the cost under Profile #2 and increases the cost under Profile #1, as well as the total expected cost.

The additional ground delay increases the cost under Profile #1 and decreases the cost under Profile #2. Overall, reducing the capacity of the arrival queue increases the total expected cost by as much as 371%, from 106.2 units for $W_{qt}^{MAX} = 56$ to 393.9 units for $W_{qt}^{MAX} = 16$.

Figure 4-10 Ground and airborne delay as a function of W_{qt}^{MAX}

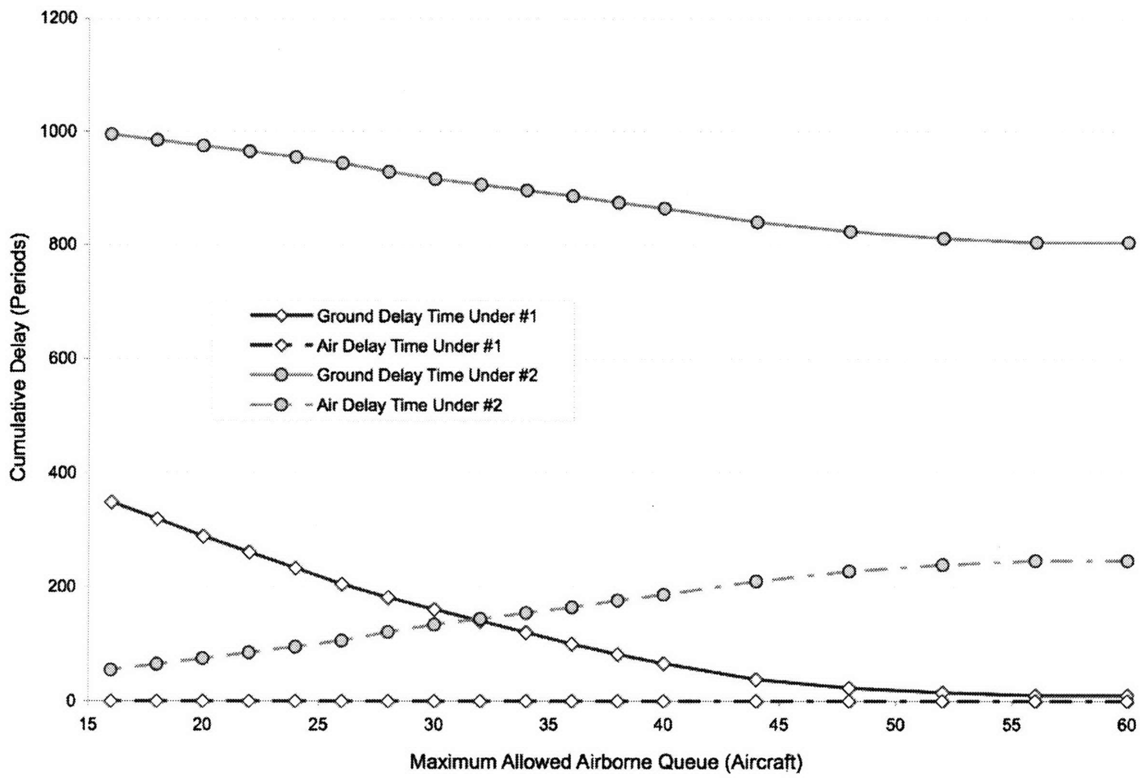

Conclusions

This experiment varies the capacity of the airborne arrival queue. If the constraint on the capacity of the arrival queue is binding in an optimal solution, decreasing the value of W_{qt}^{MAX}:

- Reduces the cost of Profile #2, which would otherwise have an arrival queue in excess of W_{qt}^{MAX}
- Increases the cost of Profile #1, which does not exhibit any airborne delays
- Increases the total expected cost of the optimal solution substantially

The infeasibility of the ECM for $W_{qt}^{MAX} < 16$ indicates that it is not possible to design a program that restricts the size of the arrival queue to an acceptable length because of the number of aircraft that are either airborne at 1300 Z or otherwise exempt. The relationship between W_{qt}^{MAX} and the time at which the decision to implement a GDP is made will be discussed further in §4.2.5.

Section 4.2.4 The Initiation Time of a GDP

The fourth experiment varies the time at which a GDP is initiated between 1300 and 1500 Z. In practice, traffic managers may postpone the decision to implement a GDP until more information is available, which is also referred to as a "wait and see" strategy. However, flights depart over time (Figure 4-11), which reduces the ability of a program to preempt airborne delays. The ECM can model the effects of the "wait and see" strategy by selectively exempting[21] all aircraft that are scheduled to depart before a

Figure 4-11 Exempt Aircraft by Decision Time

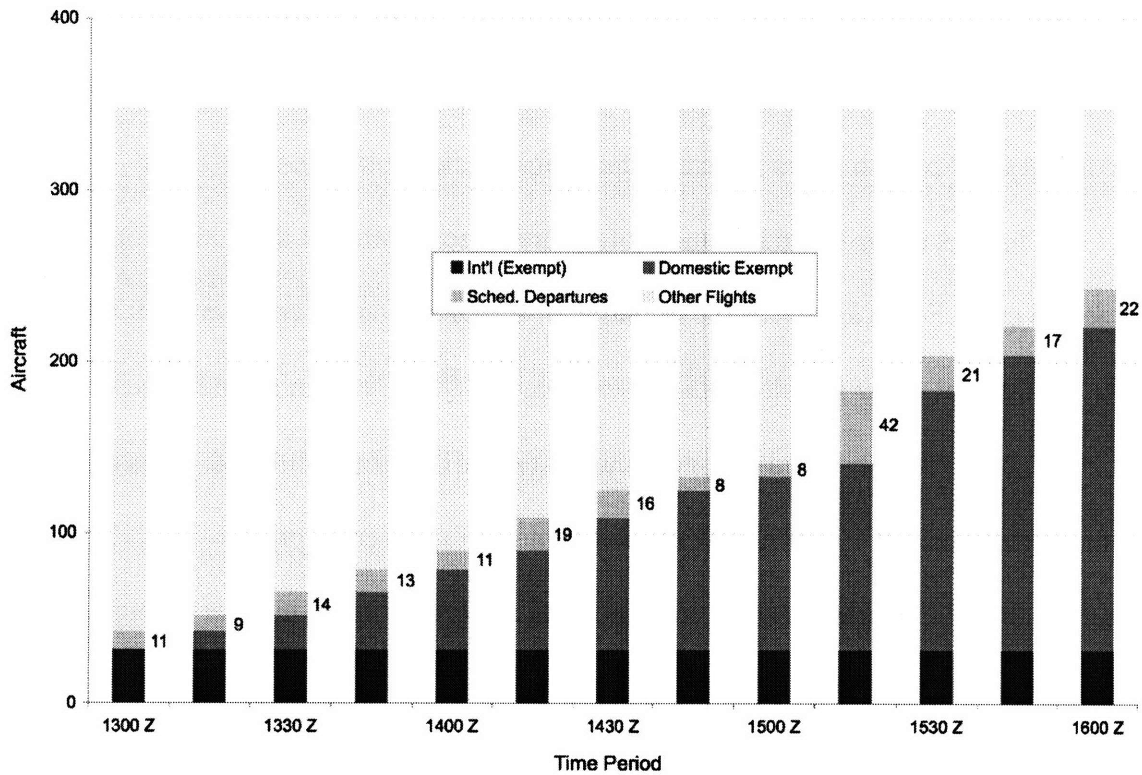

[21] §3.3.1

desired decision time. For this experiment, the time of the revision (1500 Z) is unchanged and the capacity of the arrival queue is sufficient so as to not constrain the optimal solution (Figure 4-12).

Figure 4-12 A summary of experimental setup #4

	Initial Capacity	Likelihood
Profile #1	25	0.90
Profile #2	0	0.10

	Airborne	Ground	Cumulative
Incremental Delay Cost	α_t	γ_t	χ_t

	Initiation	Revision
Decision Time	***Varies***	1500 Z

Arrival Queue Capacity	**348 aircraft**

Results

The optimal assignment of ground delay for a GDP initiated at 1300 Z results in a total expected cost of 200.4 units[22]. As the GDP is postponed, the total expected cost increases, with the greatest cost (205.1 units) incurred for an initiation at 1500 Z. However, the relative magnitude of the cost increase over this time is small: postponing the initial GDP by two hours increases the total expected cost by a mere 2.3%!

Changes under the individual profiles are more significant than those for the total expected cost. For example, as the initiation of a GDP is postponed from 1300 Z to 1500 Z, the cost under Profile #1 increases from 1,779 to 1,961 units, an increase of 10.2%. The largest observed change occurs between 1400 Z and 1415 Z. This coincides with the departure of 19 aircraft, which is the most of any period prior to 1500 Z. This suggests that the increase in costs correspond to a decrease in the number of flights that are eligible to receive ground delay.

[22] Note that this solution is equivalent to the No-MQC model in Figure 3-21, which also ignores the capacity of the arrival queue.

Conclusions

In general, the earlier that a GDP is initiated, the lower the total expected cost will be. However, for the given example, the total expected cost associated with an optimal assignment of ground delay is relatively insensitive to changes in the initiation time. Although this result may be specific to the example in this experiment, it may also signify that the total expected cost of a dynamic GDP is less sensitive to the time at which the GDP is initiated than it is to other input parameters. Postponing the initiation time reduces the number of flights that are eligible to receive ground delay to those with the latest departure times. However, a quality of an optimal solution is to delay the flights that depart last. Therefore, as long as the number of flights that remain on the ground is sufficient to avoid significant airborne delays, the "wait and see" strategy may not have significant effects on the total expected cost of the optimal GDP. The question of how many flights are sufficient may depend on other parameters, such as the capacity of the arrival queue, which was not considered in this experiment. The experiment discussed in the next section simultaneously varies both the initiation time and the capacity of the airborne arrival queue.

Section 4.2.5 The Airborne Queue Capacity and GDP Initiation Time

The fifth experiment varies the capacity of the airborne arrival queue and the timing of the initial assignment of delay simultaneously (Figure 4-13, page 102). As shown in the previous experiment, if the capacity of the arrival queue is ignored, postponing the initial assignment until 1500 Z results in a relatively small change to the total expected cost. However, a GDP initiated at 1500 Z would also result in a maximum observed arrival queue of 60 aircraft, which might exceed acceptable limits.

Figure 4-13 A summary of experimental setup #5

	Initial Capacity	Likelihood
Profile #1	25	0.90
Profile #2	0	0.10

	Initiation	Revision
Decision Time	*Varies*	1500 Z

	Airborne	Ground	Cumulative
Incremental Delay Cost	α_t	γ_t	χ_t

Arrival Queue Capacity	*Varies*

Results

This experiment consists of 69 trials, for decision times between 1300 and 1500 Z and arrival queue capacities between 16 and 60 aircraft. For eight of the trials, the ECM is infeasible (marked as "INF" in Figure 4-14), indicating that the aircraft that are already exempt at the time the GDP is initiated will exceed the capacity of the arrival queue. As a result, if the ECM is infeasible for a given trial, then a reduction in the capacity of the queue for the same decision time will also result in an infeasible solution. The minimum feasible value of W_{qt}^{MAX} for each initiation time is identified by direct experimentation and the infeasible regions of the table are shaded in light gray.

There are two manners of interpreting Figure 4-14. First, by reading across each row, the cost of the "wait and see" strategy for a desired airborne queue capacity can be identified. For example, for an arrival queue capacity of 60 aircraft, the total expected cost of the optimal solution does not change significantly if the time a GDP is initiated is postponed from 1300 to 1500 Z[23]. However, for a capacity of 30 aircraft, the formulation is infeasible unless the GDP is initiated by 1315 Z. In general, postponing the time at which a GDP is initiated does not significantly increase the total expected cost of the

[23] Note that these values are identical to those in the previous experiment for which the capacity constraint is relaxed.

optimal solution, but may preclude the possibility of reducing the size of the arrival queue to an acceptable length.

Second, reading down the columns of Figure 4-14 indicates the cost of meeting each arrival queue capacity for a given initiation time. For example, a GDP initiated at 1300 Z with an arrival queue capacity of 60 aircraft will have a total expected cost of 200.4 units. If the capacity is only 16 aircraft, then the cost is 438.6 units. In this example, the total expected cost is more sensitive to changes in the arrival queue capacity than to the time the GDP is initiated.

Figure 4-14 Total expected cost by initiation time and arrival queue capacity

	Time								
Arrival Queue Capacity	1300 Z	1315 Z	1330 Z	1345 Z	1400 Z	1415 Z	1430 Z	1445 Z	1500 Z
60	200.4	200.6	200.8	201.6	201.8	202.4	203.4	203.9	205.1
59								203.9	INF
58							203.4	INF	
56	200.4	200.6	200.8	201.6	201.8	202.4	INF		
52	201.4	201.6	201.8	202.6	202.8	INF			
48	204.8	205.4	205.6	206.4	206.6				
44	211.6	212.4	212.8	213.6	214.8				
40	230.3	231.3	232.3	233.3	INF				
39				239.3					
38	242.3	243.3	244.3	INF					
36	254.3	255.3	256.7						
34	268.1	269.1	271.3						
32	282.3	283.3	INF						
30	296.5	297.9							
28	311.5	312.9							
26	328.5	329.9							
24	349.6	351.0							
23		361.7							
22	371.0	INF							
20	392.4								
18	415.6								
16	438.8								

Units Arrival Queue Capacity: aircraft
Time: HHMM Z
Values: total expected cost for the optimal solution to the ECM
"INF" indicates an infeasible solution

Conclusions

This experiment indicates that the "wait and see" strategy may result in the maximum observed arrival queue length exceeding an acceptable limit. Furthermore, the latest time at which a program could be initiated without exceeding a given queue capacity may not be immediately clear from the optimal solution for a single trial. Alternatively, by solving the ECM at various initiation times, the time at which a decision must be made can be identified.

Section 4.2.6 Time of GDP Revision: t_{PI}

Recall that for the base arrival capacity scenario tree, ground delays are initially assigned at 1300 Z and revised once the actual arrival capacity of the airport is known with certainty. Previous experiments have assumed that the GDP revision will occur at 1500 Z. However, the time at which the revision will occur may also vary. For example, an improved forecast might be made available at an earlier time or the GDP might not be instantaneously revised once the update does become available. As the actual arrival capacity rate of the airport is known with certainty when the revision is made, this time is also referred to as the *time of perfect information*, or t_{PI}.

For the sixth experiment, t_{PI} varies between 1300 and 1600 Z (Figure 4-15). Note that the arrival capacity profiles used for each trial are the same; only the time at which the information becomes available to the traffic manager is changed. Thus, the trial for 1600 Z implies that the revision would occur one hour after the capacity reduction is first observed at the airport.

Figure 4-15 A summary of experimental setup #6

	Initial Capacity	Likelihood
Profile #1	25	0.90
Profile #2	0	0.10

Decision Time	Initiation	Revision
	1300 Z	***Varies***

	Airborne	Ground	Cumulative
Incremental Delay Cost	α_t	γ_t	χ_t

Arrival Queue Capacity	30 aircraft

Results

The total expected cost increases with each postponement of the revision time. For example, a revision at 1300 Z results in a total expected cost of 133.0 units, while a revision at 1500 Z in 196.5 units[24]. The increase in cost (123%) is much greater than that observed for a similar change in the initiation time. Furthermore, a revision after 1500 Z results in additional cost increases of 38% each at 1500 and 1515 Z.

The costs under each profile also increase with t_{PI} for most trials (Figure 4-16, page 106). Under Profile #1, cost increases most significantly (by 173%) between 1500 Z and 1530 Z, which coincides with the increase in total expected cost noted previously. Under Profile #2, the greatest observed increase in cost occurs at 1315 Z, which is also a critical time to assign *initial* ground delay[25].

However, between 1500 Z and 1530 Z, the cost of Profile #2 decreases slightly, from 1,408 to 1,401 units. The cause of this curious result is not immediately clear from the solution to the ECM. However, it suggests that, for each trial, there exist multiple optimal solutions with identical total expected cost, but for which the costs under the individual profiles are different.

[24] These results are consistent with the PI and ECM results for the example shown in §3.3.3.
[25] §4.2.4

Figure 4-16 Delay cost as a function of the GDP revision time

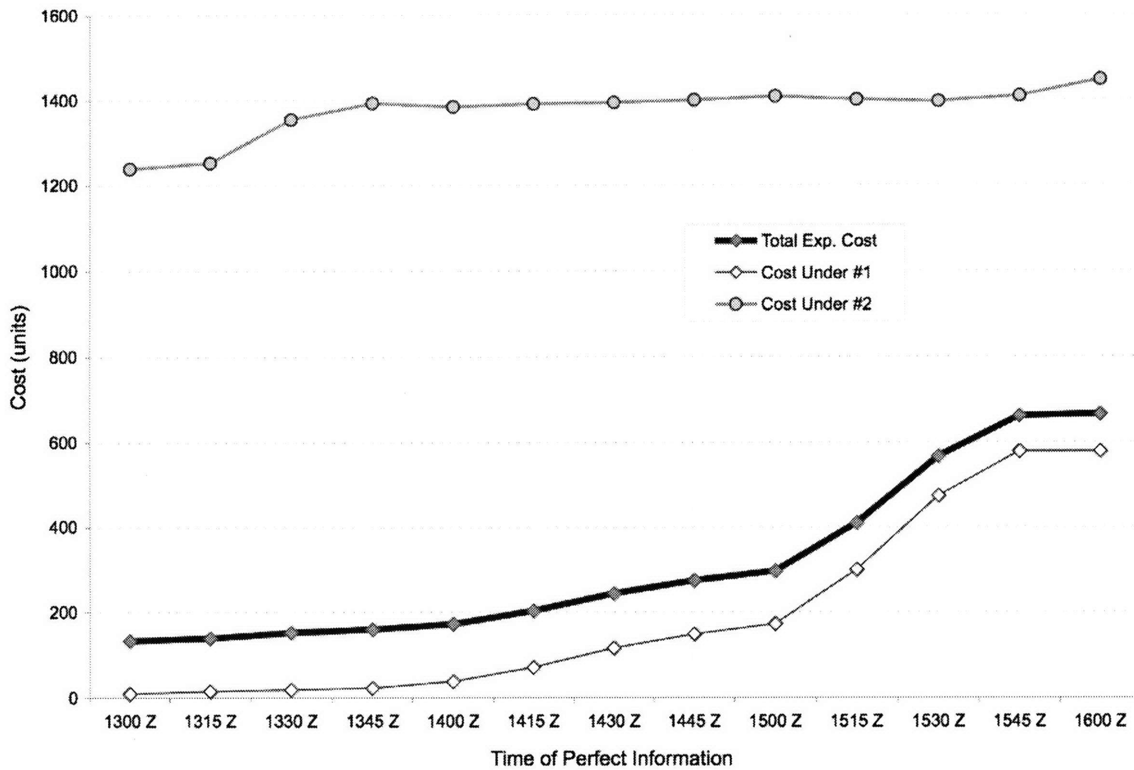

Conclusions

This experiment varies the time at which a GDP is revised to reflect the actual airport arrival capacity. Not only does the total expected cost of the optimal solution increase with t_{PI}, but the results suggest that it is much more sensitive to changes in the revision time than to the initiation time. In general, the later the revision time, the greater the total expected cost of delays.

For this example, the analysis identifies three critical decision times (1415 Z, 1500 Z, and 1515 Z), after which the delay cost increases significantly if the GDP is not revised. As the time of the revision also reflects the time at which the actual arrival capacity becomes known to the traffic manager, additional research in the sensitivity of the solution to a change in revision time would be useful for evaluating the value of technology that would increase the lead-time of capacity forecasts.

Section 4.2.7 The Maximum Airborne Delay Cost Increment: α_t'

The seventh experiment varies the incremental cost of airborne delay. The ECM considers incremental airborne delay costs that vary with the duration of delay, such as those that might be due to a flight diversion. The experiments that have been discussed previously use a hypothetical airborne delay cost function α_t[26]. In this function, the first period of airborne delay has a cost of one unit, while the fourth period of delay has an *incremental* cost of 11 units.

This experiment assumes that the incremental airborne delay cost function will follow a similar form as α_t, but that the maximum possible increment is equal to a constant A. Let the experimental incremental airborne delay cost α_t' be defined as a function of A and α_t:

$$\alpha_t' = \begin{cases} \max(A,\alpha_t) & \forall \ t \in \{1,2,3\} \\ A & \forall \ t \geq 4 \end{cases} \tag{4.01}$$

Figure 4-17 (page 108) shows a superposition of the set of cost functions that will be applied in this experiment, for which A varies between 1 and 55 units. As shown, for delays of less than three periods (45 minutes), α_t' is capped by α_t.

[26] §3.3.1

Figure 4-17 The set of airborne delay costs used for experiment #7

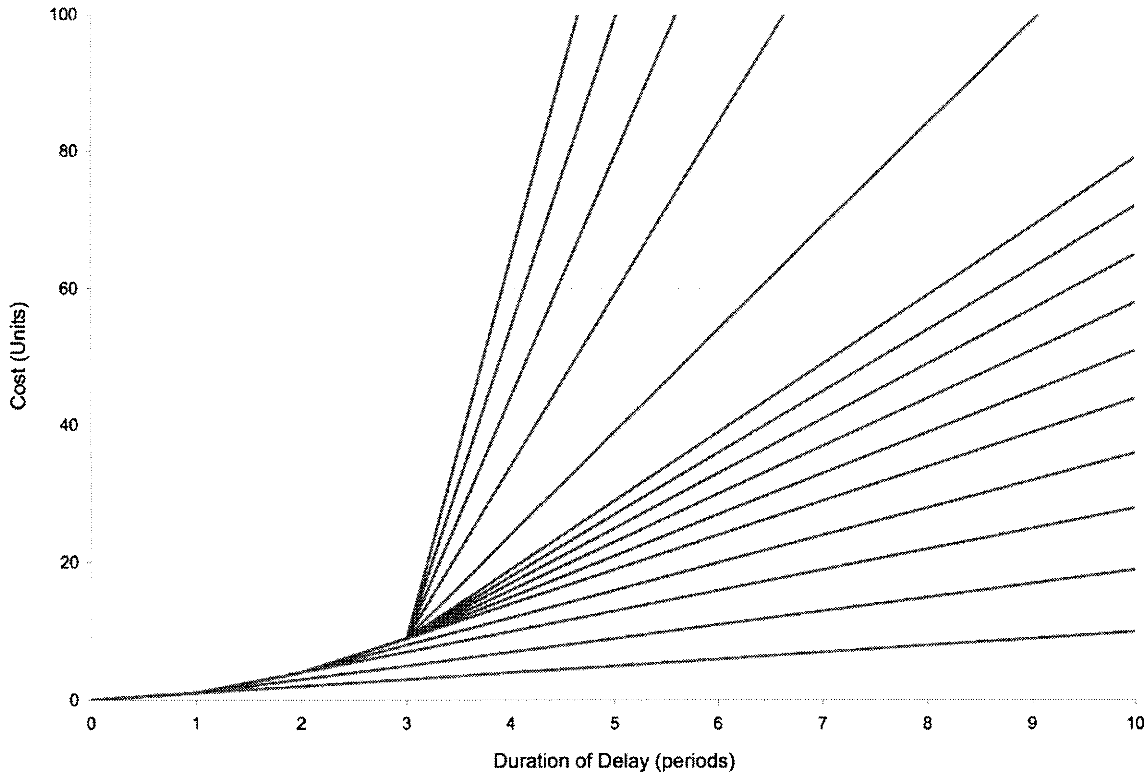

Figure 4-18 A summary of experimental setup #7 – 9

	Initial Capacity	Likelihood
Profile #1	25	**0.75**
Profile #2	0	**0.25**

	Airborne	Ground	Cumulative
Incremental Delay Cost	***Varies***	γ_t	χ_t

	Initiation	Revision
Decision Time	1300 Z	1500 Z

Arrival Queue Capacity	**348 aircraft**

In order to increase the sensitivity of the formulation to changes in A, two additional modifications are made to the experimental setup for this analysis (Figure 4-18):

- The airborne arrival queue capacity is made sufficiently large so as to not constrain the optimal solution
- The profile likelihoods are set to $\{p_1 = 0.75, p_2 = 0.25\}$

Results and Conclusions

Figure 4-19 shows the results for this experiment. As A increases, so does the total expected cost; this result is expected because A is the incremental cost of airborne delay.

In contrast to cost, the amount of ground and airborne delay assigned in the optimal solution does not always change with A. The amount of ground delay assigned under both profiles increases for $A < 5$. However, increasing A beyond 5 units does not change the assignment of delay. Furthermore, even for large values of A, the size of the maximum observed queue is never less than 44 aircraft. Previous experiments (#3 and #5) show that there exist solutions in which the observed queue length is as small as 16 aircraft. However, these solutions are not be found by increasing the value of A, alone.

Figure 4-19 A summary of the results for experiment #7

		Profile #1 (0.75)				Profile #2 (0.25)				
A	Exp. Cost	Cost	Air Delay	Ground Delay	Max. Queue	Cost	Air Delay	Ground Delay	Max. Queue	Run Time
1.00	331.5	12	0	12	0	1,290	240	810	56	4.03
2.00	368.5	17	0	17	0	1,423	230	820	55	4.37
3.00	401.5	60	0	60	0	1,426	178	872	46	4.44
4.00	411.5	65	0	65	0	1,451	172	878	45	5.17
5.00	421.0	71	0	71	0	1,471	165	885	44	5.22
6.00	422.8	71	0	71	0	1,478	165	885	44	6.04
7.00	424.5	71	0	71	0	1,485	165	885	44	6.20
8.00	426.3	71	0	71	0	1,492	165	885	44	6.25
9.00	428.0	71	0	71	0	1,499	165	885	44	6.49
10.00	429.8	71	0	71	0	1,506	165	885	44	6.12
15.00	438.5	71	0	71	0	1,541	165	885	44	6.65
25.00	456.0	71	0	71	0	1,611	165	885	44	7.22
35.00	473.5	71	0	71	0	1,681	165	885	44	7.12
45.00	491.0	71	0	71	0	1,751	165	885	44	6.95
55.00	508.5	71	0	71	0	1,821	165	885	44	7.10

Units Max. Queue: aircraft Run Time: seconds
A, Exp. Cost, Cost: units of cost Air, Ground Delay: aircraft-periods

Section 4.2.8 Airborne Delay Cost Multiplier α_t''

As in the previous experiment, the eighth experiment also varies the cost of airborne delay. However, in this case, a constant B multiplies the incremental cost of airborne delay. Let the experimental incremental airborne delay cost be defined as:

$$\alpha_t'' = B \times \alpha_t \qquad \forall \quad t \in \tau \tag{4.02}$$

Figure 4-20 shows a superposition of the incremental delay cost curves that correspond to values of B between 0 and 10. A key difference between B and A in the previous experiment is that the cost of a single period of airborne delay may exceed one unit. In all other regards, the experimental setup is the same as in the previous experiment (Figure 4-18, page 108).

Figure 4-20 The set of airborne delay costs used for experiment #8

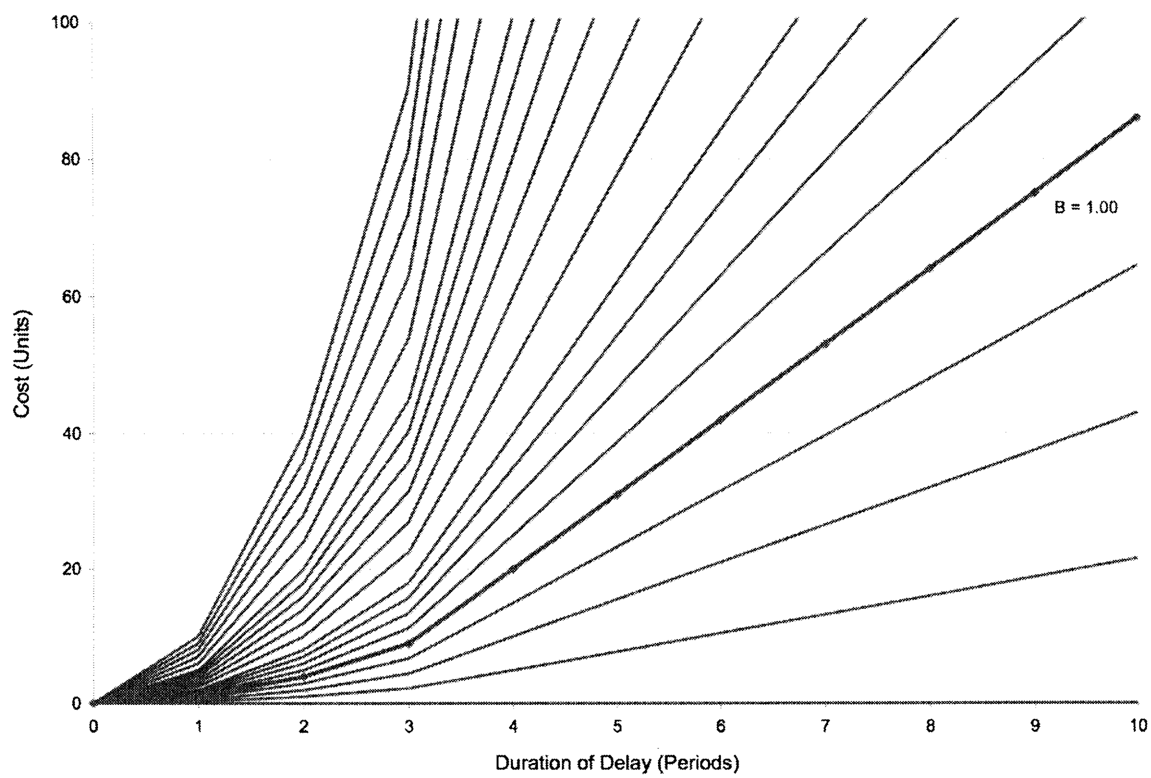

Results

The results show that increasing the multiplier B increases the total expected cost (Figure 4-21, page 112). In general, the costs under each profile also increase with B, although for several trials the cost under Profile #2 decrease. As with the experiment related to the revision time in §4.2.6, this may indicate the presence of multiple optimal solutions for these trials.

Increasing B also results in less airborne delay and more ground delay for higher values of B. The marginal benefit of increasing B, in terms of reduced airborne delay, decreases at higher values; this can be seen as a flattening of the slope of the delay curves in Figure 4-22 (page 112). Additionally, increasing B also reduces the maximum observed airborne queue size. For values of $B \geq 3.5$, the maximum observed queue is less than 30 aircraft.

As a final observation, the run time for the trial corresponding to $B = 1.50$ is 82.56 seconds, greater than any other trial in the preceding experiments. For this trial, the increased run time is due to the application of the branch-and-bound algorithm; the linear relaxation of the ECM did not give an integer solution[27] [28]. Integer solutions in the context of the ECM will be discussed further in §4.3.2.

---

[27] Repeated trials for $B = 1.50$ consistently resulted in similar model run times.
[28] The branch and bound algorithm terminated with an exact optimal solution.

Figure 4-21 Costs as a function of the multiplier B

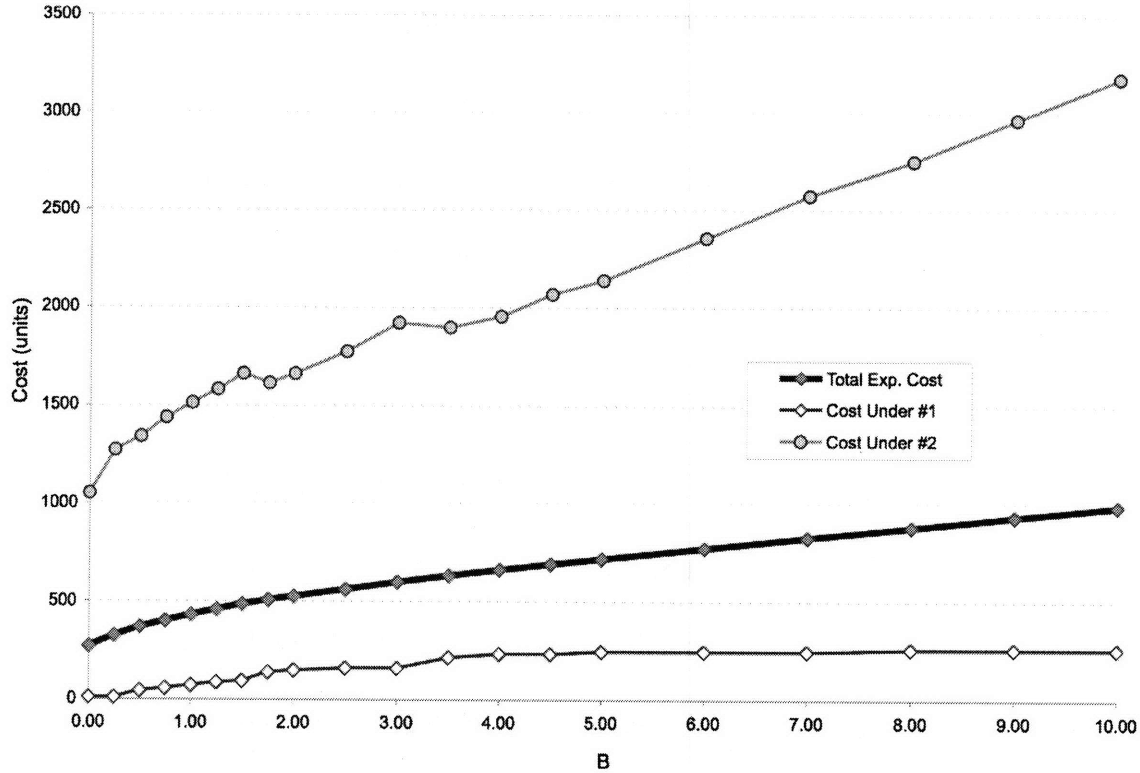

Figure 4-22 Delays as a function of the multiplier B

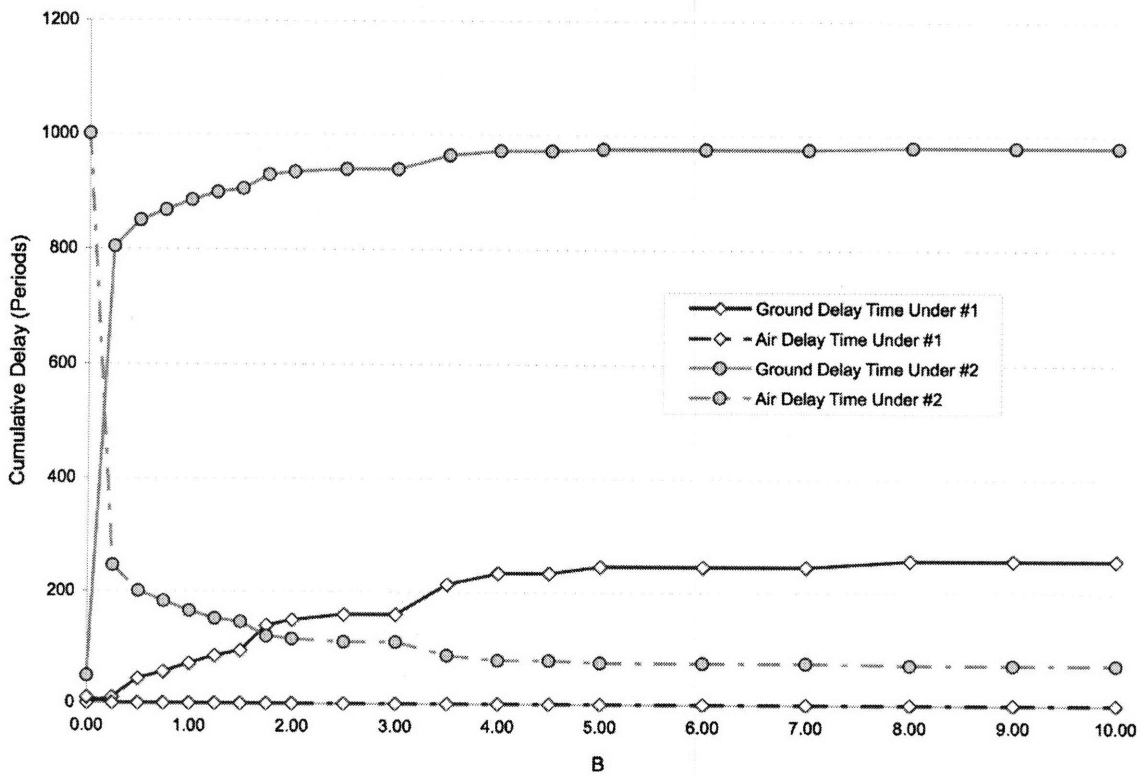

Conclusions

In this experiment, a multiplier of the airborne delay cost function is varied over a range of 0 to 10. Three general conclusions can be drawn from this experiment:

1. Total expected cost increases as airborne delay cost increases
2. The maximum observed arrival queue can be reduced by varying the cost of short-duration airborne delays
3. For one trial, the linear relaxation of the ECM resulted in a non-integer solution

The following section will discuss a third experiment of the sensitivity of the solution to changes in the airborne delays cost.

Section 4.2.9 Base Airborne Delay Cost α_t'''

As in the previous two experiments, the ninth experiment also varies the incremental airborne delay cost. In this case, a parameter C adds to the each increment such that:

$$\alpha_t''' = \alpha_t + C \qquad \forall \ t \in \tau \qquad (4.03)$$

Figure 4-23 (page 114) shows a superposition of the experimental incremental airborne delay cost functions corresponding to values of C between 0 and 25 units. Other than the cost of airborne delay, this experiment uses the same setup as the two previous experiments (Figure 4-18, page 108).

Figure 4-23 A superposition of the incremental cost curves in experiment #9

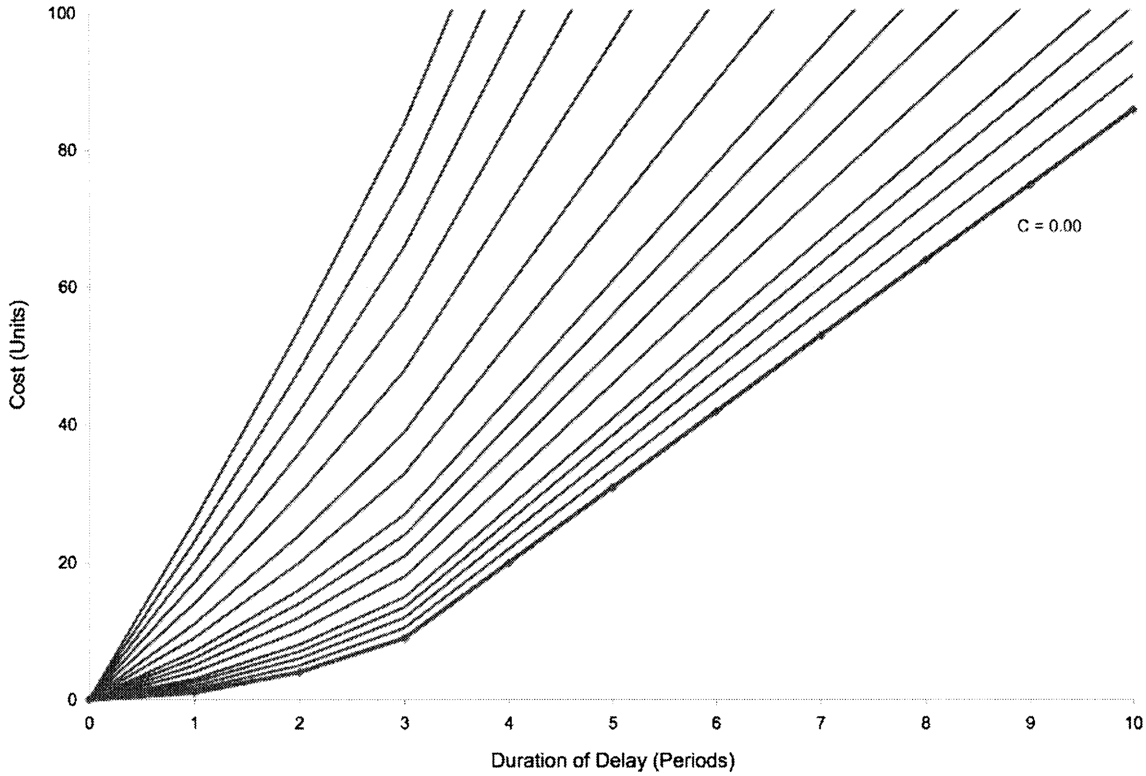

Results

For this experiment, a total of 16 trials were conducted, for values of C between 0 and 25 units. Notable values of C include:

- 0.0, for which the results are identical to that of the trial for $B = 1$
- 5.0, which is the smallest value of C (among those tested) for which the maximum observed arrival capacity is less than or equal to 30 aircraft
- 19.0, which is the smallest value of C (among those tested) for which the maximum observed arrival queue is 16 aircraft[29]

[29] The smallest value of the maximum arrival queue length for which the ECM is feasible is 16 aircraft (§4.2.3)

As with previous experiments, increasing the cost of airborne delay (in this case, by increasing C) assigns more preemptive ground delay and increases the total expected delay cost. Furthermore, for those trials with values of C greater than 16 units, the maximum observed arrival queue is reduced to the minimum feasible length. For this experiment, the linear relaxation of the ECM again failed to find an integer optimal solution in one trial (C = 1.50 units). Coincidentally, the solution in this case is the same as that discussed previously (B = 1.50).

Conclusions

The sensitivity of the solution to changes in incremental airborne delay cost was examined in three separate experiments. In the first, the maximum incremental cost varies by a parameter A (4.01); in the second, the incremental costs are multiplied by a parameter B (4.02); and in the third, a parameter C is added to each incremental cost (4.03). In each experiment, increasing the value of the parameter results in the assignment of more ground delay and an increase in the total expected delay cost. (Strictly speaking, assigned ground delay and total expected cost are *non-decreasing* with increases in the parameters A, B, and C.) Furthermore, the greater the value of the parameter in each experiment, the less the observed effect of an increase on the solution.

A key difference in the three experiments is in regards to the relative cost of short (1 or 2 periods) vs. long (> 2 periods) delays. For example, for values of $A \geq 5$ in the first experiment, further increases in A do not change the optimal solution because the incremental cost of a short delay is bounded by α_t. However, increasing the value of B

or C increases the cost of a short delay above that in the base setup – and does reduce the maximum observed arrival queue.

Section 4.3 Conclusions of the Analysis

The sensitivity analysis conducted in this chapter consisted of nine separate experiments, with a total of 195 trials. Each of the nine experiments examined the change in the optimal assignment of delay due to a change in one or more input parameters, including the arrival capacity scenario, the capacity of the airborne arrival queue, and the relative cost of airborne delay. This section summarizes the results of the experiments and identifies areas for future research.

Section 4.3.1 General Results

In practice, a GDP is an assignment of ground delay to aircraft. Each of the experiments in this analysis assigns ground delay to flights assuming a common scenario tree that contains two profiles, one representing the nominal arrival capacity and the second a reduction in capacity. An example in §3.3.2 shows that the optimal solution assigns a total of 10 periods of ground delay when the nominal capacity profile is considered alone. However, when there is also the possibility of a severe reduction, the optimal solution may assign additional ground delay under the nominal profile in order to reduce airborne delays if the reduction is realized.

As shown by the experiments in this chapter, the volume of ground delay that is assigned depends on the input parameters to the ECM. Therefore, the impact of a change in an input parameter can be measured by observing the corresponding change in ground delay assigned under the nominal profile. In each of the nine experiments, an equal or greater volume of ground delay is assigned under Profile #1 when:

- The severity of the capacity reduction increases (Experiment #1)
- The likelihood of the capacity reduction increases (#2)
- The capacity of the airborne arrival queue decreases (#3, 5)
- The GDP is initiated at an earlier time (#4, 5)
- The GDP is revised at a later time (#6)
- The cost of airborne delay increases relative to ground delay (#7, 8, 9)

In six of the nine experiments, when an increase in ground delay under Profile #1 is observed, the ground delay under Profile #2 and the total expected cost are also observed to increase. However, in three experiments – #4, 5, and 6, which correspond to the initiation and revision times – this observed relationship between total expected cost and the optimal assignment of ground delay does not hold.

First, the results of experiments #4 and #5 show that the earlier a GDP is initiated, the greater the ground delay under Profile #1 and the less the total expected cost. When the decision to assign delay is made at a later time, flights will depart and become exempt from ground delay, risking more costly airborne delays. These results suggest that the intentional postponement of the enactment of a GDP – the wait-and-see strategy – will never decrease the total expected cost of delay because it relinquishes the ability to assign delays without gaining any benefit, such as additional information.

Second, the result of experiment #6 indicates that postponing the revision time increases the ground delay assigned under Profile #1 but *decreases* the ground delay assigned under Profile #2. By postponing the revision time, the ECM applies coupling constraints to a greater number of decisions. A revision at 1300 Z will effectively allow ground delay to be assigned independently under each profile. However, a revision at 1600 Z requires that the same ground delay be assigned during the first three hours under both profiles. Therefore, as the revision time is postponed, the solutions under each profile converge, reducing the ground delay under Profile #2 and increasing the ground delay under Profile #1.

Section 4.3.2 Run Time of the ECM

For the 195 trials conducted as part of the sensitivity analysis, there were 187 feasible solutions. The cumulative distribution of run times for these 187 trials is shown in Figure 4-24.

For the 187 unique trials, the average run time was 7.58 seconds and the maximum 89.61 seconds. For three of the trials (1.5%), the linear relaxation of the ECM did not result in an integer solution and the branch-and-bound algorithm was applied[30]. Excluding these three trials, the run times of the model range from 3.38 to 12.42 seconds, with an average of 6.69 seconds. As shown, the distribution of run times is also observed to vary across the different experiments (Figure 4-25). A more detailed exploration of the relationship between the variation of different input parameters and model run time is left as an area for future research.

[30] For each of these, the branch-and-bound algorithm terminated with an optimal integer solution.

Figure 4-24 Cumulative distribute of run times

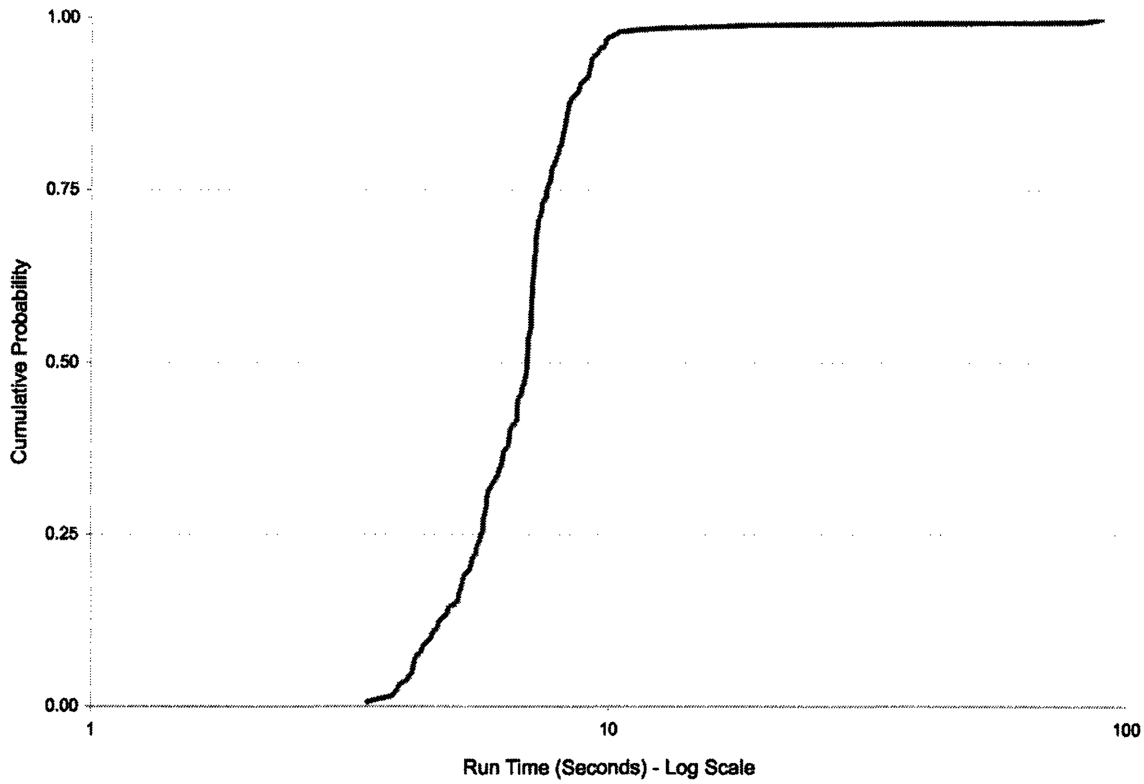

Figure 4-25 Distribution of model run times by experiment

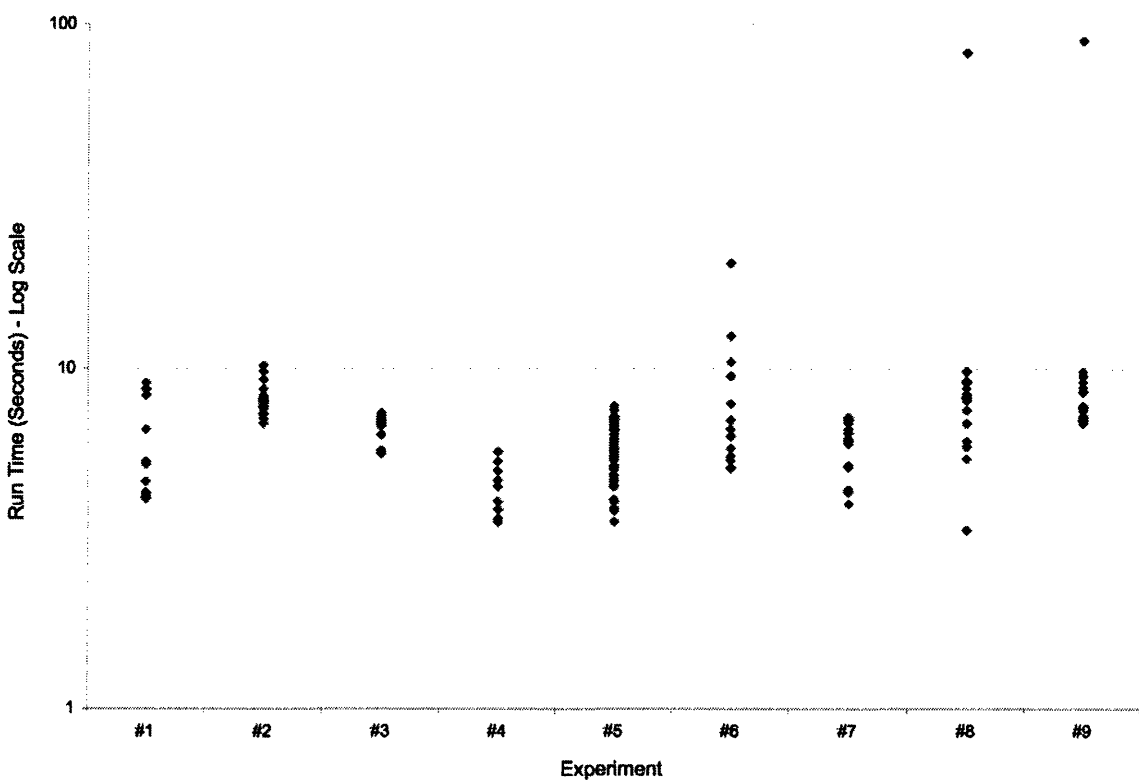

As a final note regarding the run time of the ECM, recall that the comparison of the ECM to prior models in the literature (§3.3.4) indicated that the ECM requires significantly longer to solve than the other models. However, the sensitivity analysis also demonstrates that the ECM solves a hypothetical problem quickly enough to be of practical use – even when the linear relaxation does not yield an integer solution. The analysis also shows that model run times tend to be relatively consistent for different values of the input parameters; in only three of the nearly 200 trials did the solution require longer than 12 seconds.

Section 4.3.3 Sensitivity in Practice

The sensitivity analysis presented in this chapter assumes a hypothetical airport arrival capacity scenario. However, the results of the analysis offer three meaningful conclusions for how a similar analysis could be used for the design of ground delay programs in practice.

First, a sensitivity analysis identifies the effects of uncertainty in the value of input parameters upon the optimal assignment of ground delay. For example, there may exist several different methods to forecast the arrival capacity of an airport, each of which may provide a different estimate of the likelihood of a decrease in capacity. Instead of choosing a single value, a sensitivity analysis could explore the impact on the assignment of ground delay for a range of possible probabilities.

Second, a sensitivity analysis demonstrates the effect of a wait-and-see strategy. For the hypothetical example used in this discussion, postponing the initial decision until 1500 Z resulted in only a small relative increase (2.3%) in the total expected cost.

However, if the decision is made after 1315 Z, it is not possible to prevent the airborne arrival queue from exceeding 30 aircraft. By simultaneously varying both the capacity of the arrival queue and the time of the initial decision, it is possible to identify critical times by which a decision should be made in order to achieve a desired arrival queue capacity.

Third, a sensitivity analysis demonstrates the value of improving the accuracy and precision of various input parameters. As shown in experiment #6, revising the GDP at an earlier time may significantly reduce the total expected cost. However, the time of the revision also depends on the time at which a new arrival capacity forecast is made. For a single GDP, a sensitivity analysis could be used to determine the reduction in total expected delay cost corresponding to an earlier revision. Further research could extend this analysis in order to demonstrate the practical value of developing new technology that delivers improved forecasts with greater lead-time.

Section 4.3.4 Summary

The nine experiments discussed in this chapter examine the sensitivity of the optimal solution to changes in the input parameters and demonstrate three key results. First, the ECM is capable of solving for the optimal assignment of ground delay under a wide range of input parameters and could be used to perform a sensitivity analysis of the decision to implement a GDP. Second, the time required for the ECM to find a solution for various combinations of inputs is short enough that the model could be used during the design of a GDP in practice. Third, the analysis demonstrates how the optimal assignment of delay varies with respect to changes in different input parameters. For example, “critical” decision times are identified by varying the time of the initial and revising decisions.

Furthermore, this analysis also illustrates the opportunity for future study on the effects of variations in the value of the input parameters on the assignment of ground delay. Specifically, studies could be conducted to explore the sensitivity of the solution to simultaneous variation in multiple variables, or to determine the value of improving forecasting technology. A third area for exploration is to examine how the ECM assigns delay under various arrival capacity scenarios – an analysis that will be presented in the next chapter.

Chapter 5 An Application of the ECM to Hypothetical Arrival Capacity Scenarios

The ECM is a mixed-integer program that solves the single airport ground hold problem (SAGHP) subject to a stochastic and dynamic forecast of the airport arrival capacity. This model improves upon previous stochastic/dynamic SAGHP models in the literature because it considers both non-linear airborne delay costs and the capacity of the airborne arrival queue.

Chapter Five presents an analysis comparing the solution to the ECM to those of other models in the SAGHP literature for ten different hypothetical airport arrival capacity forecasts. The various forecasts, which are referred to as "case studies," include uncertainty about the severity and timing of a reduction in the arrival capacity rate, as well as the times at which revised forecasts are made available. The analysis shows that not only can the ECM be applied to a wide variety of arrival capacity forecasts, but it also results in solutions that are preferable to those of previous models in the literature.

The analysis in Chapter Five is presented in three parts. Section 5.1 discusses the objectives of the analysis and describes the common experimental setup. Section 5.2 presents the results of 10 arrival capacity cases. Section 5.3 summarizes the results and discusses the merits of the ECM.

Section 5.1 An Overview of the Analysis

The analysis is presented as a set of ten experiments, each of which compares the solution to the ECM to those of other models for a given arrival capacity forecast. Each forecast is represented as a scenario tree that contains a set of hypothetical arrival capacity profiles and one or more forecast revision times. The set of scenario trees is chosen to include uncertainty in various elements of an arrival capacity forecast, such as the number of profiles and the start time, severity, and duration of a decrease in arrival capacity. Each scenario tree, together with the other inputs that are considered by the ECM, is referred to as a "case study."

The case study analysis makes three general contributions to this thesis. First, it demonstrates how arrival capacity scenario trees might represent various hypothetical forecasts, including those for which the likelihoods of individual profiles change over time, and demonstrates that the ECM can solve for the optimal assignment of ground delay given these forecasts. Second, it shows that the solutions of the ECM are preferable to those of previous models in the SAGHP literature for various forecasts of the arrival capacity. Third, it shows that the ECM could also be used as an off-line tool to identify strategies for the design of a GDP subject to stochastic and dynamic arrival capacity forecasts.

Section 5.1.1 Experimental Setup

Each of the ten cases is based upon a common experimental setup, previously introduced in §3.3.1, which includes a set of expected flight arrivals, a capacity of the airborne arrival queue, and delay cost functions. Arrival demand is given by a list of 348 aircraft that are scheduled to arrive at ORD over a four-hour period of time[31]. The capacity of the arrival queue is assumed to be 30 aircraft for all time periods. And the functions used to evaluate the cost of a solution are given by the incremental costs in Figure 3-9.

The arrival capacity scenario tree is different for each of the ten cases in this analysis. However, the differences reflect various elements of an arrival capacity forecast and are also chosen so as to facilitate the comparison of results between each experiment. For example, Case #2 represents a forecast with an uncertain time of an increase in capacity, while Case #5 represents uncertainty in the time of the *decrease*. Furthermore, four cases demonstrate how the timing of a forecast revision changes the optimal assignment of ground delay (Figure 5-1).

Figure 5-1 A description of the ten arrival capacity case studies

Case #	Type(s) of Uncertainty Represented	Forecast Revisions
1	*Severity* of a decrease in capacity (2 alternatives)	(2) 1400, 1500 Z
2	*Duration* of a decrease in arrival capacity	(1) 1530 Z
3	*Duration* of a decrease in arrival capacity	(4) 1530, 1600, 1630, 1700 Z
4	*Severity* of a decrease in capacity (5 alternatives)	(1) 1500 Z
5	*Start time* for a decrease in capacity	(1) 1500 Z
6	*Start time* for a decrease in capacity	(4) 1500, 1530, 1600, 1630 Z
7	*Start time* and *severity* of a decrease in capacity	(1) 1500 Z
8	*Start time* and *severity* of a decrease in capacity	(3) 1500, 1600, 1700 Z
9	*Start time* and *duration* of a decrease in capacity	(1) 1500 Z
10	*Start time* and *duration* of a decrease in capacity	(3) 1500, 1600, 1700 Z

[31] As shown in Figure 3-07

Models Compared

Each experiment compares the solution to the ECM to those of various other SAGHP models for a given case study. Three of these models are from the SAGHP literature[32]:

1. The Mukherjee fully dynamic model (M-DM)
2. The Mukherjee fully dynamic model with distributed delays (M-DMDD)
3. The Richetta-Odoni partially dynamic model (RO-PDM)

An additional three models are adapted from the ECM in order to test various assumptions.

4. The capacity of the arrival queue is sufficiently large so as to accommodate all arrival demand (No-MQC)
5. The arrival capacity profile, as it will be realized, is revealed at 1300 Z (PI)
6. All flights are exempt from ground delay (No-GDP)

Basis for Comparison

The various models are compared on the basis of the total ground and airborne delay times, the maximum size of the observed airborne arrival queue, and the total expected delay cost that would be incurred by each solution. As the models may have different objective functions, the term "cost" always refers to the value obtained by applying a set of common cost functions (Figure 3-9) to a given solution. These functions express the delay costs accrued by each flight due to ground, airborne, and cumulative delay.

[32] For more information on these models, please refer to the discussion in §3.3.1.

Section 5.1.2 Terminology

To facilitate a discussion of the results, several terms are defined here; each refers to the solution to a given model. A complete set of definitions is provided in the appendix.

- Cost under a profile: the cost of the solution to a model as evaluated by the set of common cost functions in Figure 3-9 and assuming the realization of a particular arrival capacity profile
- Total expected cost: the sum of the costs under each profile weighted by the profile likelihoods
- Maximum observed arrival queue: the maximum number of flights that would simultaneously experience airborne delay under any profile
- Maximum observed flight delay: the greatest amount of delay assigned to any one flight under any profile
- Departure queue: At a given point in time, the number of aircraft being delayed by a GDP past their scheduled departure time

Section 5.2 Case Studies on the Arrival Capacity Scenario

This section presents the results of the ten experiments based upon the arrival capacity case studies outlined in Figure 5-1 (page 125). Each experiment solves for the assignment of ground delay for a given case study using the various SAGHP models and the solutions are discussed and compared.

Section 5.2.1 Case #1: Revising the Profile Likelihoods

The first case study is an adaptation of the example previously presented in §3.3.3. The scenario tree in this case assumes that there are three possible arrival capacity profiles with likelihoods $\{p_1 = 0.8, p_2 = 0.1, p_3 = 0.1\}$. Under each of the first two profiles, the airport arrival capacity is the nominal rate of 25 aircraft/period. Under the third profile, the arrival capacity is reduced to 0 aircraft/period between 1500 and 1600 Z and then increases to the nominal rate thereafter. In a practical sense, this scenario is identical to that of the example shown previously in §3.3.3. Although the initial scenario at 1300 Z contains three profiles instead of two, the likelihood of a nominal arrival capacity rate being realized is still 0.9.

As in the original example, the scenario tree in Case #1 assumes that a GDP may be initiated at 1300 Z and then revised at 1500 Z to reflect the actual arrival rate capacity of the airport. However, a key difference in this case is that the arrival capacity forecast will also be revised at 1400 Z to indicate whether or not Profile #1 will be realized (Figure 5-2). The additional forecast has the practical effect of revising the likelihood of each profile. If the revised forecast indicates that Profile #1 will be realized, then the forecast follows the upper branch of the scenario tree, for which the relative probability of a reduction in capacity is 0. On the other hand, if Profile #1 is ruled out as a possibility at 1400 Z, then the relative likelihood of Profile #2 becomes 0.5.

Figure 5-2 Scenario tree for Case #1

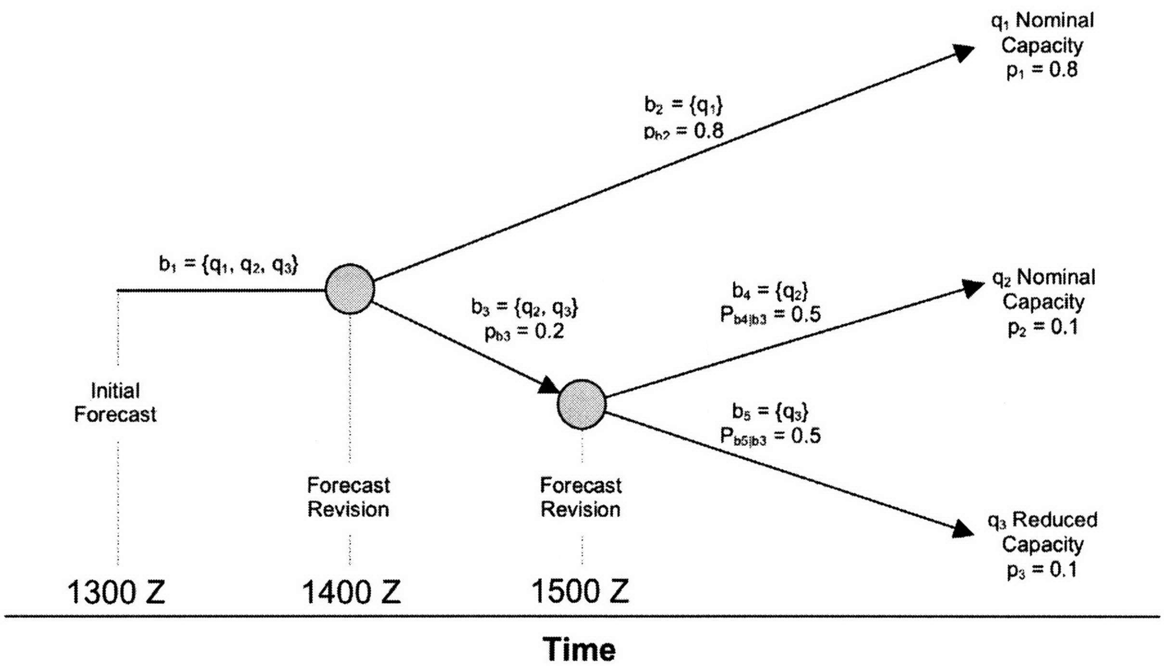

Figure 5-3 Results for Case #1

	ECM	M-DMDD	M-DM	RO-PDM	PI	No-MQC	NoGDP
Exp. Total Cost	187.7	188.8	187.7	196.3	133.0	181.3	490
% of ECM Cost		100.6%	100.0%	104.6%	70.9%	96.6%	260.8%
Max Obs. Queue	30	52	46	44	16	40	92
Exp. Delay							
Ground	141.5	97.2	99.0	96.1	109.2	110.3	0.0
Airborne	11.4	19.8	18.9	21.4	4.8	16.3	114.0
Max Flight Delay							
Airborne	6	7	11	15	5	7	0
Ground	4	4	4	4	4	4	4

Units	Exp. Total Cost:	units of cost	Exp. Delay:	aircraft-periods
	Max Obs. Queue:	aircraft	Max Flight Delay:	periods

Results

For Case #1, the total expected cost of the optimal solution to the ECM is 187.7 units and the maximum observed airborne arrival queue is 30 aircraft (Figure 5-3). In comparison to the other models from the literature, the solution to the ECM has the lowest cost and the smallest maximum observed arrival queue length. The solution to M-DM is second-most preferable, with a cost that is identical to the ECM (M-DMDD is 0.6% higher) but a significantly longer maximum observed arrival queue length.

As a result of the additional forecast revision at 1400 Z, the total expected delay cost of the solution to the ECM is reduced by 37% as compared to that of the original example in §3.3.3. Furthermore, the realization of Profile #1 (in Case #1) results in a delay cost of 38 units, which is a 78% reduction as compared to the nominal profile if there is no revision (173 units). However, under Profiles #2 and #3 the cost is approximately equivalent to what it would be without the revised forecast.

The effect of the 1400 Z revision is also shown in the graph of the departure queue over time (Figure 5-4). The optimal solution assigns significantly less ground delay under Profile #1 than it would if there were no revision but similar amounts of ground delay under Profiles #2 and #3. Thus, adding a forecast revision at 1400 Z reduces the total delay cost if Profile #1 occurs and has a minimal effect if Profile #1 is not realized.

This case study makes two important points. First, it demonstrates how a scenario tree could represent an arrival capacity forecast in which the likelihood of an event changes over time. Second, it shows that the improvement of the ECM over other models depends on the timing and availability of information. For the example in §3.3.3, which has the same initial forecast and also reveals the true capacity at 1500 Z, the total expected cost of the solution to the ECM is much greater than that of other models. However, for Case #1, which adds an additional revision at 1400 Z, the solution to the ECM is equal to or lower in cost than those of other models from the literature.

Figure 5-4 Ground queue size by time for the solution to the ECM to Case #1

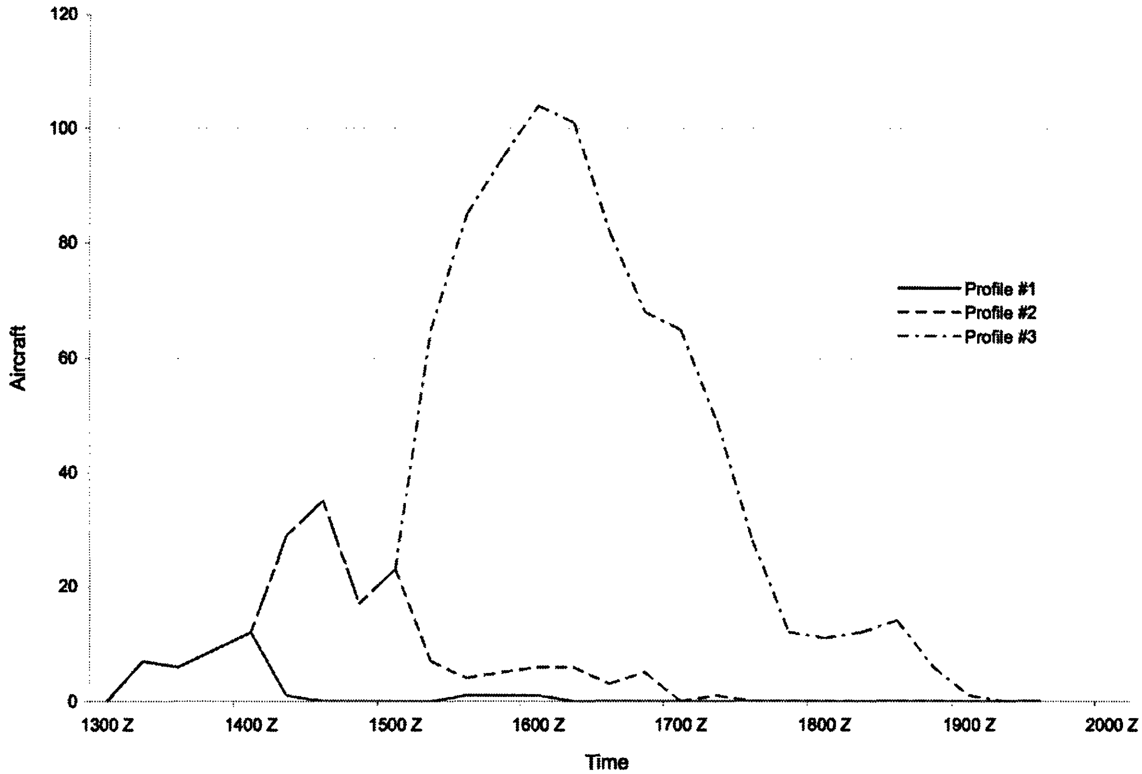

Section 5.2.2 Case #2: An Increase in the Arrival Capacity

The scenario tree in Case #2 represents a hypothetical forecast in which the time of an increase in the arrival capacity rate is uncertain. For this scenario, the initial arrival capacity of the airport is assumed to be 10 arrivals/period, which might represent a reduction from the nominal capacity due to local weather conditions, such as a high cross wind or poor visibility. It is assumed that, at some time in the future, the conditions causing the reduction will abate and the arrival capacity rate will increase to the nominal level of 25 arrivals/period. The time at which the increase will occur is uncertain.

An initial forecast of the arrival capacity at 1300 Z predicts that the airport arrival capacity will increase at one of five possible times, each with a given likelihood. Which profile will be realized will be indicated by a revised forecast at 1530 Z.

1. Capacity increases at 1530 Z ($p_1 = 0.10$)
2. Capacity increases at 1600 Z ($p_2 = 0.20$)
3. Capacity increases at 1630 Z ($p_3 = 0.30$)
4. Capacity increases at 1700 Z ($p_4 = 0.20$)
5. Capacity increases at 1730 Z ($p_5 = 0.20$)

Figure 5-5 and Figure 5-6 show the arrival capacity profiles and scenario tree.

Results

The total expected cost of the optimal solution to the ECM for Case #2 is 882.0 units (Figure 5-7, page 134). Of the other models, the cost of M-DMDD is equal to that of the ECM and those of M-DM and RO-PDM are 4.5% and 27.8% higher, respectively. All solutions, except that for RO-PDM, avoid the possibility of airborne delay.

The results of Case #2 suggest a strategy for reducing the cost of delays in practice. The optimal solution to the ECM proactively assigns ground delay immediately prior to the forecast revision. In the half hour preceding the revision, between 1500 and 1530 Z, the departure queue increases from 19 to 72 aircraft (Figure 5-8, page 134). Then, at 1530 Z, ground delays are revised according to the profile that will be realized. For example, if Profile #1 (the earliest possible increase) is realized, all aircraft in the departure queue are immediately released. However, if the increase does not occur until 1730 Z, then additional ground delays are assigned and the departure queue builds to 141 aircraft. These results show that, in practice, it may be preferable to enact a GDP before a forecast revision, rather than after, in order to create the option to assign additional ground delay if the capacity does not increase.

Figure 5-5 Superposition of the arrival capacity profiles for Case #2

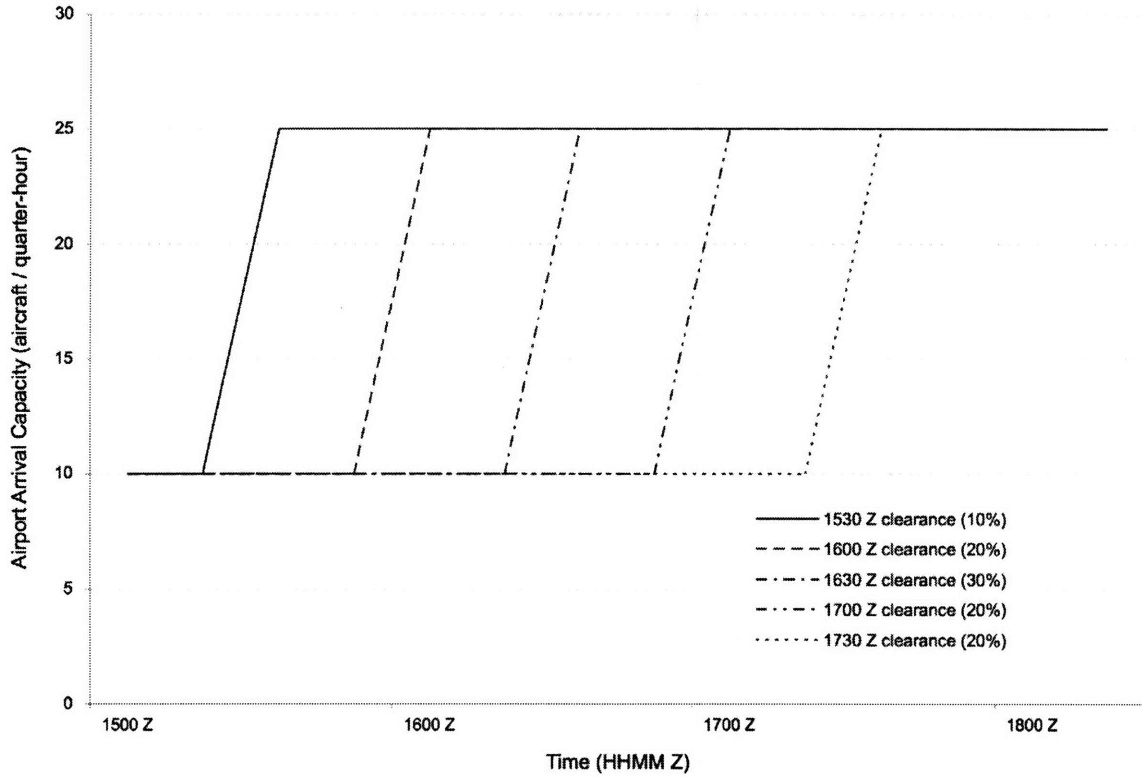

Figure 5-6 Arrival capacity scenario tree for Case #2

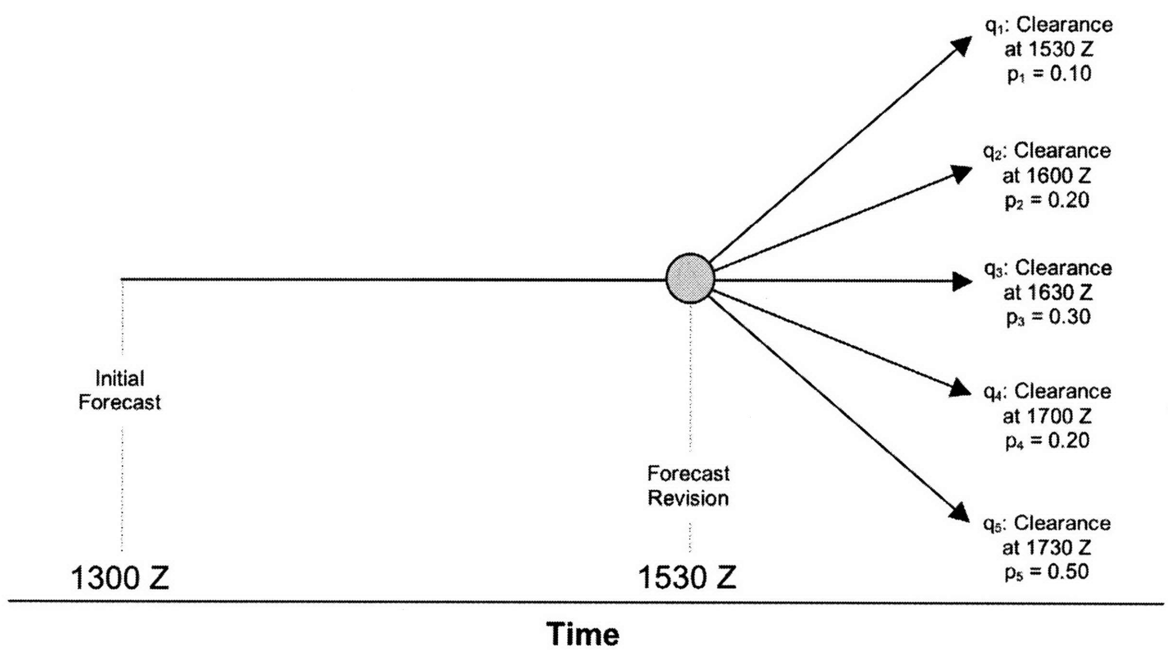

Figure 5-7 Results for Case #2

	ECM	M-DMDD	M-DM	RO-PDM	PI	No-MQC	NoGDP
Exp. Total Cost	882.0	882.0	921.7	1,127.0	859.5	882.0	4,265
% of ECM Cost		100.0%	104.5%	127.8%	97.4%	100.0%	483.6%
Max Obs. Queue	0	0	0	33	0	0	130
Exp. Delay							
Ground	882.0	882.0	882.0	886.2	859.5	882.0	0.0
Airborne	0.0	0.0	0.0	49.4	0.0	0.0	859.5
Max Flight Delay							
Airborne	8	8	12	15	7	8	0
Ground	0	0	0	3	0	0	6

Units Exp. Total Cost: units of cost, Exp. Delay: aircraft-periods
Max Obs. Queue: aircraft Max Flight Delay: periods

Figure 5-8 Ground, or Departure, Queue for Case #2, ECM

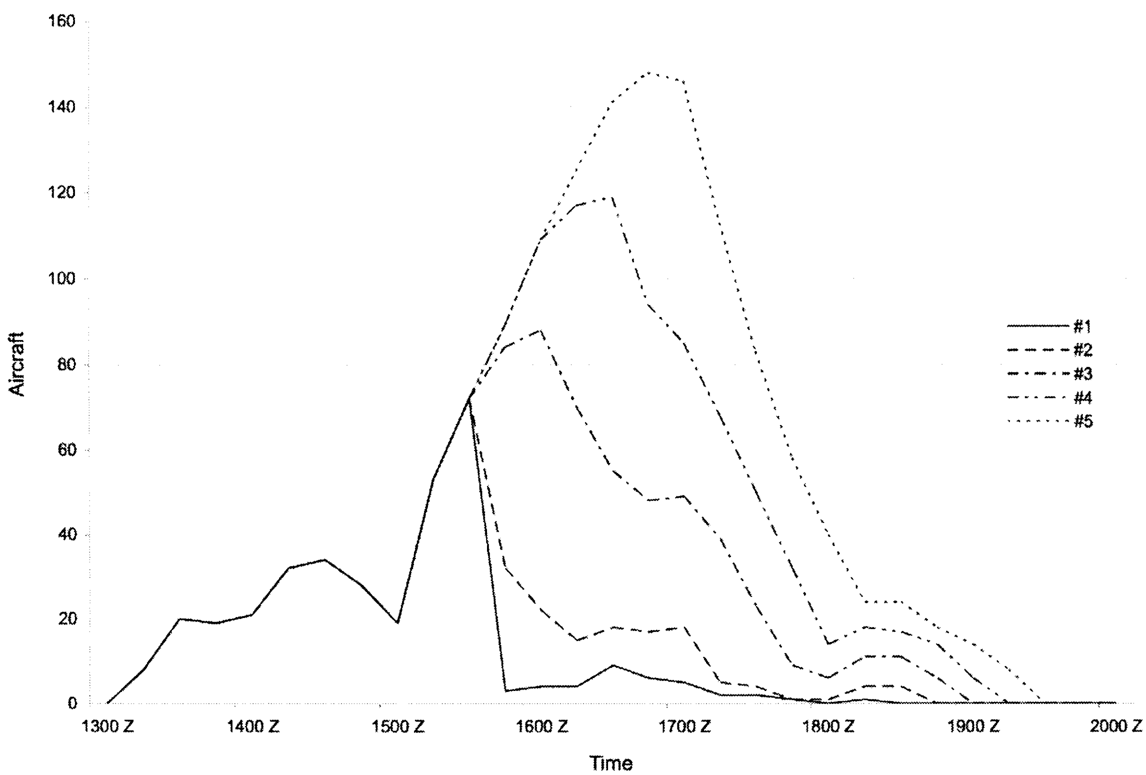

Case #2 also highlights important shortcomings in both M-DM and RO-PDM. First, the solution to M-DM for this example has a higher cost than those of other models because it assigns more lengthy ground delays to individual aircraft. For example, M-DM, M-DMDD, and the ECM each assign a total of 1,446 periods of ground delay under Profile #5. However, of these models, only the solution to M-DM assigns more than eight periods of ground delay to any one flight. In fact, M-DM assigns delay in excess of eight periods to 83 flights, with a maximum observed flight delay of 12 periods. Recall that the cost of ground delay to a flight is assumed to increase after the eighth period because of the likelihood of a cancellation[33]. However, as shown by the other solutions, these costs may be avoided in this example by either considering non-linear ground delay costs directly (as by the ECM) or by distributing ground delay more evenly (M-DMDD).

Second, RO-PDM results in solutions with a greater total expected delay cost because this model is not fully-dynamic. While the optimal strategy for the solution to the ECM requires assigning delay in anticipation of a revision at 1530 Z, RO-PDM is not able to consider solutions that would release these flights at 1530 Z if Profile #1 is realized. Instead, for the RO-PDM solution, aircraft remain in the departure queue at 1530 Z even if the arrival capacity increases (Figure 5-9, page 136). As a result, the cost due to ground delay under Profiles #1 and #2 increases substantially and the solution to RO-PDM is 27.8% higher in cost than the ECM.

[33] §3.3.1

Figure 5-9 Departure queue for Case #2, RO-PDM

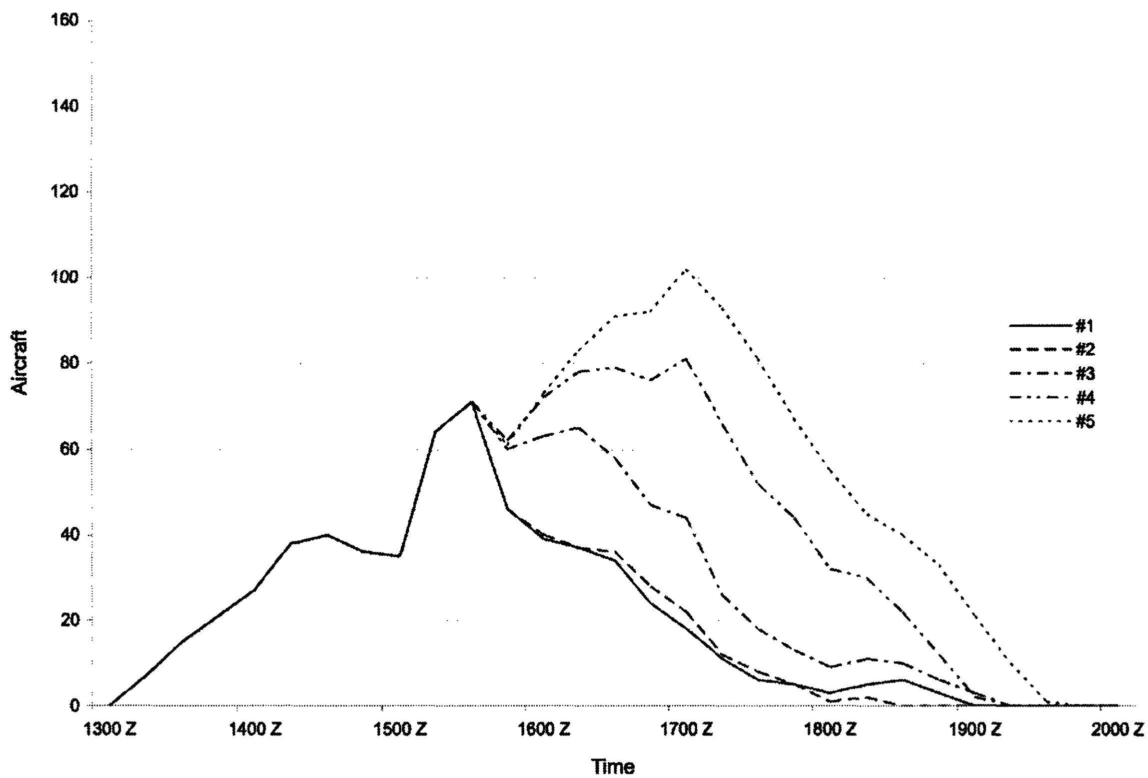

Section 5.2.3 Case #3: The Capacity Increase for an Alternative Scenario Tree

The solution to the ECM for Case #2 highlights the relevance of the 1530 Z revision to the optimal solution. Ground delays are initially assigned to the flights that are scheduled to depart during the half hour prior to the revision in order to create an option to assign additional delay, if necessary, at 1530 Z. To show that this strategy is also used when there are multiple revisions, Case #3 presents a scenario tree in which the actual arrival capacity rate is revealed gradually, over a series of additional forecasts. In this case, the revision at 1530 Z only indicates whether or not Profile #1 will be realized. If Profile #1 is not realized, then there may be up to three additional forecast revisions that will gradually reveal which arrival capacity profile will occur (Figure 5-10).

Figure 5-10 Scenario tree for Case #3

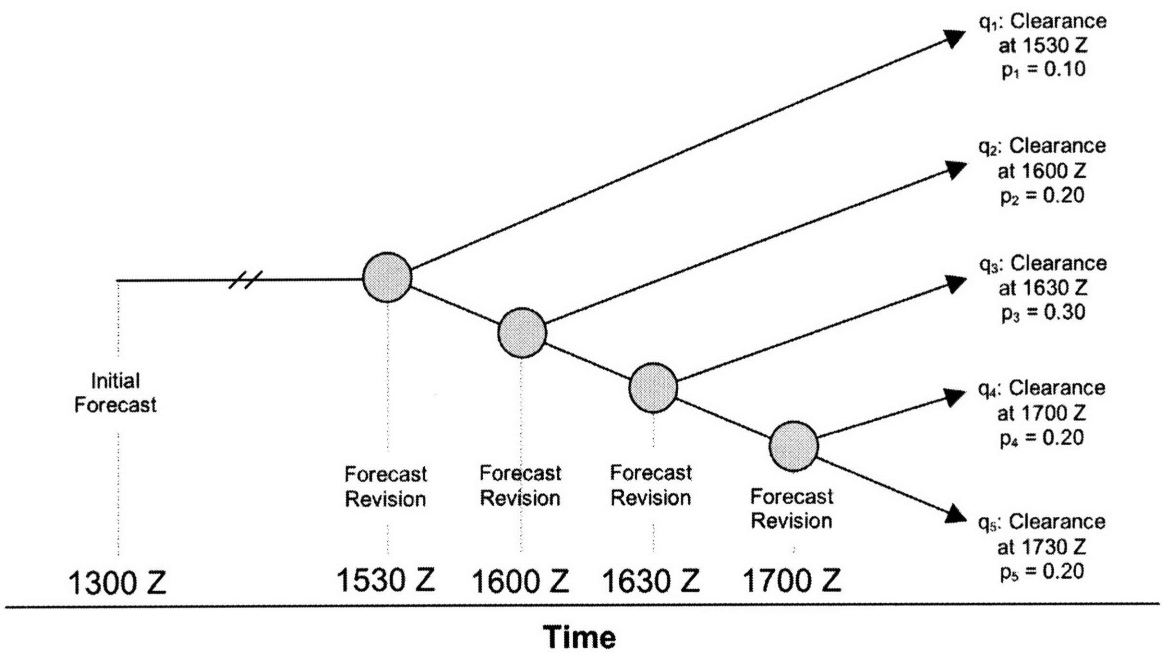

Results

For Case #3, the optimal solution to the ECM results in a total expected delay cost of 952.0 units, which is an increase of 70.0 units over the solution to Case #2. The cost increases because Case #3 is, in effect, a postponement of the time of the complete forecast revision and postponing the forecast increases the total delay cost of the optimal solution[34]. The solution to each of the different models[35] also demonstrates an increase in cost. However, the magnitude of the cost increase is not uniform across the different models. For example, while the cost of the M-DMDD solution is equal to that of the ECM for Case #2, it is greater for Case #3 by 18.1 units (Figure 5-11, page 138). The relative performance of the different models subject to changes in the revision time will be discussed further in §5.2.6.

---

[34] The sensitivity analysis in §4.2.6 explores the relationship between cost and the revision time in greater detail.

[35] except PI, for which the time of the revision is irrelevant

Figure 5-11 Results for Case #3

	ECM	M-DMDD	M-DM	RO-PDM	PI	No-MQC	NoGDP
Exp. Total Cost	952.0	970.1	1,015.1	0.0	859.5	0.0	4,265
% of ECM Cost		101.9%	106.6%	0.0%	90.3%	0.0%	448.0%
Max Obs. Queue	25	21	21	0	0	0	130
Exp. Delay							
Ground	835.4	817.1	815.1	0.0	859.5	0.0	0.0
Airborne	50.5	56.8	56.8	0.0	0.0	0.0	859.5
Max Flight Delay							
Airborne	8	10	0	0	7	0	0
Ground	2	3	3	0	0	0	6

Units	Exp. Total Cost:	units of cost	Exp. Delay:	aircraft-periods
	Max Obs. Queue:	aircraft	Max Flight Delay:	periods

Section 5.2.4 Case #4: A Reduction of Uncertain Severity

Case study #4 shows an example in which there is a possibility of an extreme outcome, the closure of an airport to flight arrivals, for an extended period of time. In this case, it is known that the arrival capacity of the airport will be reduced between 1500 and 1800 Z, with an initial forecast at 1300 Z and one revision at 1500 Z. However, the severity of the decrease is uncertain; the initial scenario contains five arrival capacity profiles, each with an equal likelihood:

1. 20 arrivals/period $(p_1 = 0.20)$
2. 15 arrivals/period $(p_1 = 0.20)$
3. 10 arrivals/period $(p_1 = 0.20)$
4. 5 arrivals/period $(p_1 = 0.20)$
5. 0 arrivals/period $(p_1 = 0.20)$

Results

For Case #4, the optimal solution to the ECM has a total expected cost of 2,768.6 units and exhibits a maximum observed arrival queue of 30 aircraft. Each of the other models results in a solution with both a higher cost and a longer maximum observed arrival queue (Figure 5-12). In particular, M-DMDD is 10.3% higher in cost (and could also result in a maximum observed arrival queue of 50 aircraft).

The solution to M-DMDD is also notable because it has a higher cost despite assigning less airborne delay and less ground delay. The cost of the M-DMDD solution is higher because the airborne delays that would occur are longer in duration. As compared to the ECM, M-DMDD assigns less proactive ground delay, which results in an airborne arrival queue that forms earlier than in the solution to the ECM. As a result, under Profile #5, 47 aircraft receive more than three periods of airborne delay in the M-DMDD solution, as compared to only 25 aircraft for the ECM.

Figure 5-12 Results for Case #4

	ECM	M-DMDD	M-DM	RO-PDM	PI	No-MQC	NoGDP
Exp. Total Cost	2,768.6	3,054.0	3,163.0	3,542.0	2,567.8	2,768.6	15,947
% of ECM Cost		110.3%	114.2%	127.9%	92.7%	100.0%	576.0%
Max Obs. Queue	30	50	48	63	30	30	260
Exp. Delay							
Ground	1,929.8	1,835.4	1,845.0	1,853.4	1,730.4	1,929.8	0.0
Airborne	153.2	112.6	107.8	136.6	152.4	153.2	1,882.8
Max Flight Delay							
Airborne	13	14	18	23	13	13	0
Ground	12	12	12	12	12	12	12

Units	Exp. Total Cost:	units of cost	Exp. Delay:	aircraft-periods
	Max Obs. Queue:	aircraft	Max Flight Delay:	periods

Section 5.2.5 Additional Case Studies

Six additional cases are presented for which the arrival capacity scenario trees represent forecasts with uncertainty in the start time, severity, and duration of a decrease in the arrival capacity. For comparison, the six cases show three separate sets of arrival capacity profiles with different assumptions as to the number of forecast revisions (Figure 5-13). Figures showing the scenario trees of the individual experiments are included in the appendix.

Figure 5-13 Additional case studies performed

Scenario Description	Count of Profiles	Forecast Revisions	
		One (at 1500 Z)	Many
Uncertain *start time* for a decrease in capacity	5	Case #5	Case #6
Uncertain *start time* and *severity* of a decrease in capacity	9	Case #7	Case #8
Uncertain *start time* and *duration* of a decrease in capacity	9	Case #9	Case #10

In each case, the solution to the ECM results in a lower total expected cost than those of M-DMDD, M-DM, or RO-PDM. For brevity, the results of these six cases are discussed as part of the conclusions in the next section.

Section 5.2.6 A Summary of the Experimental Results

For each of the ten arrival capacity case studies presented in this chapter, the ECM compares favorably to prior models in the literature. First, the solution to the ECM has an equal or lower total expected delay cost than the other models. Second, the solution to the ECM does not result in a maximum observed arrival queue of more than 30 aircraft under any profile. Although the limit of 30 aircraft was chosen arbitrarily for these examples, the ECM may be solved with any upper limit specified by the user. Based on these experimental results, two general conclusions are drawn regarding the performance of the various SAGHP models used in this experiment.

First, the experimental results show that the relative performance of the individual models is consistent across the different cases. Figure 5-14 (page 143) compares the total expected cost of each model for each case, where cost is expressed relative to the cost of the ECM. The costs of the various solutions exhibit a similar order in each case. For example, in each experiment, PI has an equal or lower total expected cost than each of the other models, while RO-PDM has a higher cost than other models.

As a general observation, the various models can be placed in an approximate hierarchy according to the total expected cost of their solutions to each case:

$$\text{PI} \leq \text{No-MQC} \leq \text{ECM} \leq \text{M-DMDD} \leq \text{M-DM} \leq \text{RO-PDM} \leq \text{No-GDP}$$

In Figure 5-14, the models are arranged in columns corresponding to this order; models with lower costs, such as PI, are on the left and those with higher costs, such as No-GDP, are on the right. In general, the ECM has a lower cost than every other model, except PI and No-MQC. There exist two notable exceptions to this hierarchy, which are highlighted in the figure.

The first exception is the total expected cost of the ECM for Case #0, which refers to the simple example that is discussed in §3.3.3. For this case, the ECM has a higher cost than the other models from the literature. The reason for the higher cost is that the ECM assigns proactive ground delay to limit the size of the arrival queue to 30 aircraft, a limit that is exceeded by M-DM, M-DMDD, and RO-PDM (Figure 5-15, page 143). If the arrival queue is not constrained, as shown by No-MQC, then the cost of the solution is lower than that of the other models. Thus, one exception to the hierarchy is cases in which the ECM results in a higher cost because it avoids arrival queues with a length that might be unacceptable in practice.

The second exception is the cost of the M-DM and M-DMDD solutions for Case #1. Although M-DMDD results in a lower cost than M-DM in most examples, in Case #1, M-DMDD has a higher cost. The only difference between these two models is the objective function of the formulation; M-DMDD adds a penalty that increases with the duration of ground delay (§3.3.1). As shown by Case #2, this cost penalty causes M-DMDD to distribute ground delay more evenly, resulting in fewer ground delays in excess of eight periods and a lower total expected cost for some cases. However, this penalty also increases the cost of ground delay relative to that of air delay, which may increase the amount of airborne delay in the solution. In Case #1, the additional airborne delay increases the total expected cost and also results in a larger maximum observed arrival queue. Thus, while M-DMDD distributes ground delay more evenly than M-DM, it may also increase the size of the arrival queue beyond acceptable limits.

The additional airborne delay that may be incurred for solutions to M-DMDD represents a significant drawback to the manner used by this model to distribute ground delay among different flights. The ECM is also able to distribute ground delay by incorporating a super-linear cumulative delay cost function, which, by itself, would also increase the cost of ground delay relative to airborne delay. However, unlike M-DMDD, the ECM avoids the potential pitfalls associated with distributing ground delay more evenly by explicitly considering both non-linear airborne delay costs and the capacity of the airborne arrival queue.

Figure 5-14 Comparison of the total expected costs for various models by case

Total Expected Cost (as a % of the ECM cost) by Case							
Case	PI	No-MQC	ECM	M-DMDD	M-DM	RO-PDM	No-GDP
#0	44.9%	67.6%	100.0%	68.3%	68.5%	70.9%	165.1%
#1	70.9%	96.6%	100.0%	100.6%	100.0%	104.6%	260.8%
#2	97.4%		100.0%	100.0%	104.5%	127.8%	483.6%
#3	90.3%		100.0%	101.9%	106.6%		448.0%
#4	92.7%	100.0%	100.0%	110.3%	114.2%	127.9%	576.0%
#5	95.7%		100.0%	100.2%	115.9%	119.1%	720.3%
#6	85.0%		100.0%	104.8%	120.0%		639.6%
#7	69.8%	88.5%	100.0%	100.7%	104.1%	112.5%	409.2%
#8	54.5%	83.8%	100.0%	100.3%	102.4%		319.5%
#9	94.4%		100.0%	100.2%	108.4%	114.9%	562.9%
#10	78.8%		100.0%	104.9%	109.1%		470.1%

Figure 5-15 Comparison of the maximum observed arrival queues by model, case

Maximum Observed Arrival Queue (in aircraft) by Case							
Case	PI	No-MQC	ECM	M-DMDD	M-DM	RO-PDM	No-GDP
#0	16	54	30	56	56	56	92
#1	16	40	30	52	46	44	92
#2	0		0	0	0	33	130
#3	0		25	21	21		130
#4	30	30	30	50	48	63	260
#5	0		16	17	17	21	150
#6	0		27	38	38		150
#7	30	65	30	98	98	108	260
#8	30	74	30	105	105		260
#9	0		16	16	16	21	140
#10	0		25	30	30		140

Note 1: The models are arranged in order of total expected cost from lowest to highest. Exceptions to this order are highlighted in gray.

Note 2: Blank cells indicate that a model was not used to solve that particular case.

A second conclusion drawn from the experimental results is that the magnitude of the benefit offered by the ECM over other models in regards to the total expected costs depends on the times at which the arrival capacity forecast is revised. For example, the arrival capacity scenario in Case #6 contains the same profiles as that in Case #5, but postpones certain revisions to the arrival capacity forecast. As a result of the postponement, the relative difference between the costs of the solutions to M-DMDD and the ECM increases from 0.2% to 4.8%. However, postponing information does not always increase the benefit of the ECM. For example, in Case #1, the ECM is equal in cost to M-DM, but for Case #0, the ECM has a higher cost. Further exploration of the relationship between the arrival capacity scenario and the relative cost of the ECM over other models is recommended as an area for future research.

Section 5.3 Conclusions

The arrival capacity case studies discussed in this chapter demonstrate that the ECM can be applied to various hypothetical arrival capacity forecasts. One example shows how a scenario tree could be used to capture a forecast in which the likelihoods change over time and others exhibit uncertainty in various elements of the arrival capacity forecast. In each of these examples, the ECM compares favorably to previous models in the literature, especially for cases in which there is a possibility of a severe reduction in arrival capacity for an extended period of time. This section summarizes the contributions of the ECM and discusses the implications of this model in regards to the SAGHP literature and the design of GDPs in practice.

Section 5.3.1 Comparison to SAGHP Models

The Extended Cost Model builds upon prior models in the SAGHP literature, such as those of Mukherjee (2004), to also consider an airborne arrival queue capacity and non-linear functions for the cost of delays. However, the formulation of the ECM requires a significant number of variables to capture non-linear delay costs[36], which results in a model run time that is greater than those of other models in the literature, including those of Mukherjee. For the largest cases discussed in this thesis, the ECM had approximately 1 million variables and 2 million constraints, and required 10 minutes to solve[37]. Thus, the run time of the ECM could be a barrier for practical implementation.

In conclusion, the ECM offers several key improvements over prior models in the SAGHP literature, including:

1. The ability to avoid solutions that result in airborne queues of unacceptable length
2. Direct consideration of non-linear cost functions
3. A more even distribution of delay without an increase in the total expected delay cost due to additional airborne delay

However, these benefits come at the expense of model size and run time.

Section 5.3.2 Thoughts on the Use of SAGHP Models in Practice

Ultimately, the value of models in the SAGHP literature will be determined by how well they meet the needs of practitioners. In contrast to the current GDP design process, which depends on one or more individuals to make subjective decisions in

---

[36] §3.3.3
[37] For platform information, please refer to §3.3.1

regards to the scope and planned arrival rates of a program, the ECM offers two main advantages. First, the ECM objectively considers both the uncertainty of arrival capacity forecasts and the dynamic nature of information. Second, the ground delays assigned by the ECM minimize the total expected delay cost, while avoiding the possibility of arrival queues of unacceptable length.

However, before the ECM may be applied to the design of GDPs in practice, there remain three significant hurdles. First and foremost, the stochastic and dynamic forecasts of airport arrival capacity that are assumed by the ECM would need to be developed in the form of scenario trees. Currently, these types of forecasts are not widely available and research that has sought to develop them has met with limited success[38]. Second, the manner in which the ECM might be used in concert with other mechanisms such as slot-credit-substitutions, compression, and swaps[39] needs to be clarified. Mukherjee (2004) offers an initial discussion of how to include participant actions and Chapter Seven of this thesis will also discuss this topic. Third, further research should try to improve the run time of the ECM. Although the ECM solves for the optimal assignment of ground delay in a reasonable amount of time for the examples discussed in this thesis (348 flights, 2-10 profiles, ~40 time periods), it is not known if these examples are as large as those that might be encountered in practice.

Despite these hurdles, the ECM still offers significant benefit in its current form for use as an analytical tool. For example, examination of the solutions to the many case studies in this thesis indicates that the optimal solution often holds flights on the ground in anticipation of a revision to the forecast. It is possible that another, faster algorithm

[38] Liu (2007), MIT Lincoln Laboratory (2004)
[39] Chang et al. (2001)

could be designed for implementing such a strategy in practice. Further research could explore how the ECM and other SAGHP models might be used to identify other strategies that might be of use in practice.

Section 5.3.3 Moving Forward

Although the ECM compares favorably to other SAGHP models, there are many important questions that remain to be addressed, such as how equity might be defined in regards to a GDP. As the ECM, as well as other SAGHP models, is motivated by the design of GDPs in practice, the further improvement of SAGHP models should also be directed by practical objectives. However, these practical objectives vary across the different groups that are affected by the outcome of GDPs. For example, traffic managers seek to ensure the efficient use of airspace, commercial air carriers desire to retain flexibility and control of their schedules, and passengers seek efficient and inexpensive service.

The next chapter in this thesis seeks to formalize the process of understanding and evaluating the objectives of a GDP. Starting with an institutional perspective, the needs of stakeholders and their relationships will be considered in conjunction with the use of GDPs. Then, in Chapter Seven, the ECM will be adapted to illustrate how GDPs might be modified to reflect the needs of some of these stakeholders.

Chapter 6 Ground Delay Programs and the Stakeholder Perspective

The FAA uses Ground Delay Programs (GDPs) to manage the arrival demand at capacity-constrained airports. By delaying flights on the ground before they depart, a GDP may reduce the volume of costly airborne delays, as well as air traffic congestion, that occurs when the number of aircraft that seek to land at an airport exceeds its arrival capacity. However, in order to delay flights on the ground, a GDP must be initiated before they depart, at a time when the future weather conditions – and the arrival capacity of the destination airport – may be uncertain.

In practice, the design of a GDP is a collaborative process between a national air traffic manager at the FAA and representatives from various commercial airlines. The traffic manager decides how much delay should be assigned and how to allocate it among flights. This decision is based upon the traffic manager's evaluation of an uncertain and dynamic forecast of the airport arrival capacity and the possible effects of a GDP on air traffic flow management throughout the National Airspace System (NAS). To form this evaluation, the traffic manager draws on a combination of personal experience and judgment and opinions solicited from the airlines and local air traffic control centers. Once delays are assigned, airlines may adjust them (subject to certain limitations) to reflect their own internal objectives. A criticism of the design in practice is that it relies heavily on personal judgment – two traffic managers might interpret the same forecast differently.

In the academic literature, the design of GDPs is represented as the single airport ground hold problem (SAGHP). SAGHP models are advantageous for use in the design of GDPs because they assign ground delays that are both consistent and optimal for a given objective function. The development of these models, such as the "Extended Cost Model" (ECM) presented in §3.2, has largely focused on capturing the objectives of the design of GDPs as considered by the traffic manager in practice. In Chapters Six and Seven, the design of a GDP is approached from a new direction. Chapter Six draws upon stakeholder theory and suggests that the design of GDPs in practice should give more consideration to the passenger. Chapter Seven adapts the ECM to consider passengers and demonstrates that the cost of delays to passengers could be significantly reduced if the traffic manager were to consider them directly during the design of a GDP.

Figure 6-1 Conceptual model of the research approach

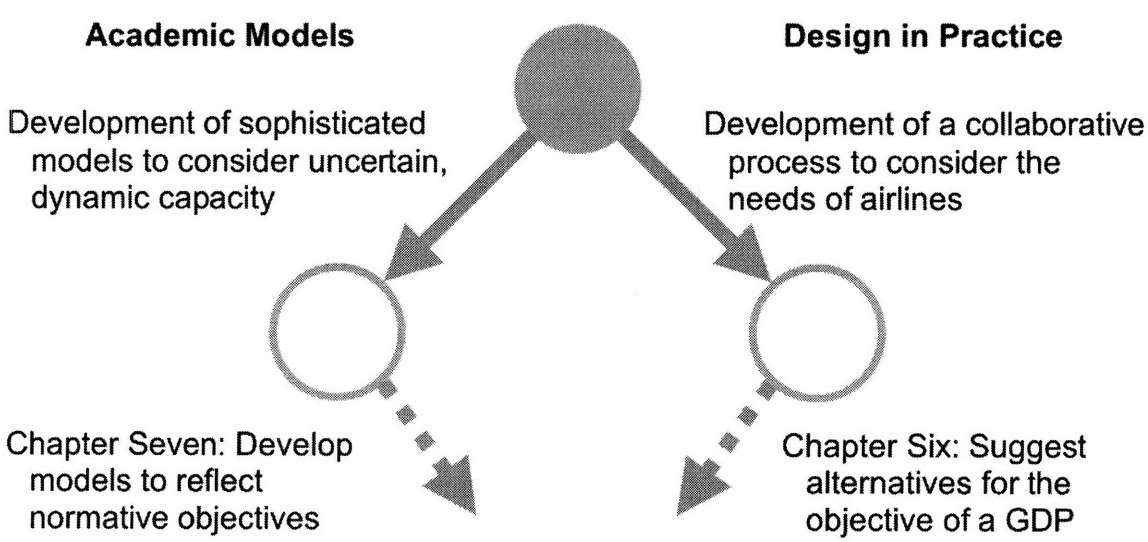

The analysis presented in Chapter Six tests the hypothesis that the design of GDPs in practice excludes consideration of groups that should receive greater representation. The analysis determines the salience of various stakeholders groups, where salience is defined as the importance of a group in regards to the design of a GDP. Section 6.1 introduces the stakeholder methodology and identifies key stakeholder groups. Section 6.2 describes the relationships between the various stakeholders. Section 6.3 examines the impacts of a GDP. And Section 6.4 determines the salience of various stakeholder groups and compares this to the stakeholder roles and impact on the design of GDPs in practice.

Section 6.1 A Stakeholder Approach to GDPs

The objective of the analysis presented in this chapter is to test the hypothesis that the design of GDPs in practice excludes consideration of groups that should receive greater representation. Existing literature related to the design of GDPs has focused on improving the design to reflect those groups that are already part of the design process. In order to identify additional groups, this analysis draws upon the body of literature that deals with stakeholder theory. The theory of stakeholders is adapted for use in thinking about the design of GDPs and an analysis is conducted to determine the salience of various stakeholders in the GDP design process. By comparing these results to the design in practice, those groups that should receive more consideration are identified.

Section 6.1.1 Stakeholder Identification

The first step of the analysis is to identify the various stakeholders and to determine the relationship between each in the context of a GDP. A GDP *stakeholder* is defined as any group that "affects the outcome of or is affected by the outcome of a Ground Delay Program." This definition is adapted from a more general definition of stakeholders as proposed by R. Edward Freeman[40]. The concept of a stakeholder that is proposed in this thesis is much broader than that suggested in existing SAGHP literature. Whereas the literature has limited the discussion of GDPs to include groups that influence the outcome of a GDP, this broader definition would also apply to groups that are affected by programs but which do not have an active part in the decision process, such as passengers.

Under this broader definition, various entities are identified, including commercial airlines, general aviation, industry trade groups, airports, and businesses that provide goods and services to passengers and airlines. Each is either an active participant in the design or implementation of a GDP, mentioned in the academic or technical literature in conjunction with GDPs, or an intermediary or third party between the FAA and one of the other groups. These entities may be organized into a taxonomy, an example of which is shown in Figure 6-2.

[40] Freeman (1984)

Figure 6-2 A taxonomy of various GDP stakeholders[41]

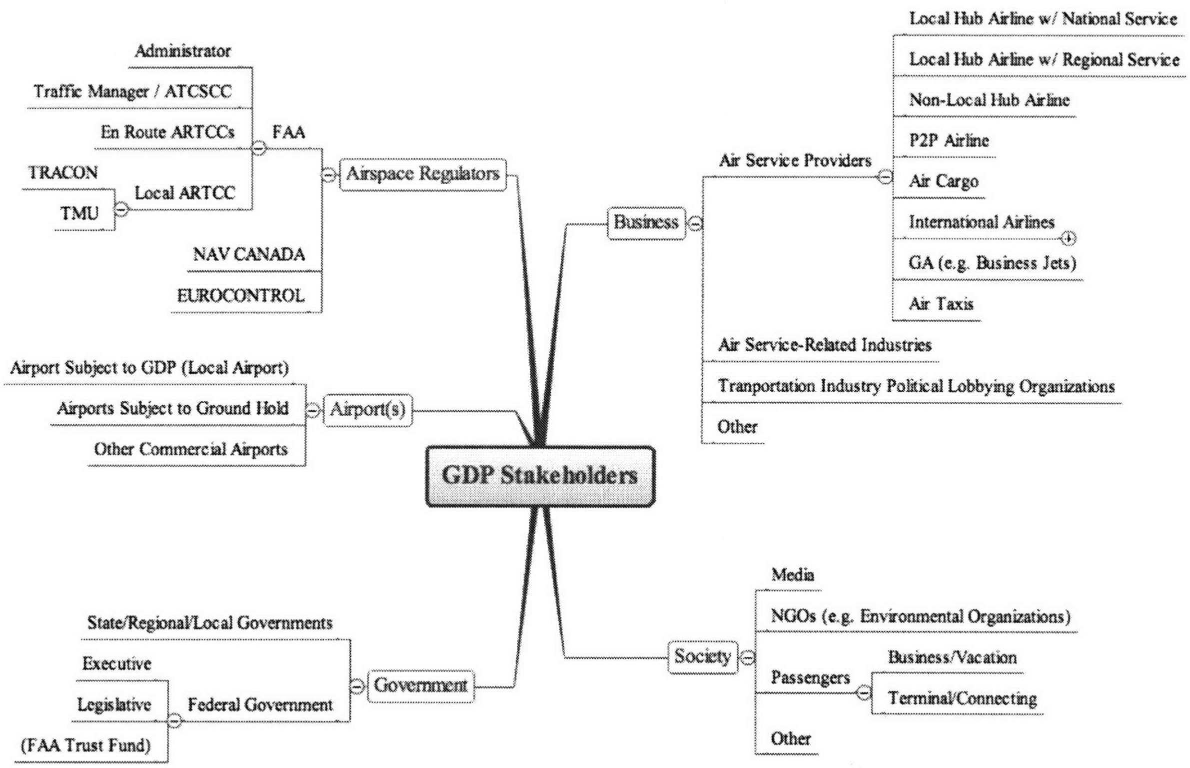

As based on the taxonomy, the various entities can be grouped into eight stakeholder groups (Figure 6-3, page 154). Each group is broadly defined so as to encompass many sub-groups. For example, Air Transportation Service Providers, or ASPs, includes large commercial carriers, air taxis, and general aviation. Furthermore, groups are also defined by the role they play in regards to a GDP. For example, passengers could be considered as part of society as a whole, but they are unique among other groups in society because passengers, themselves, are delayed by a GDP. All stakeholders identified previously are included in one of these eight groups (except those that are related to but not affected by a specific GDP, such FAA research organizations).

[41] Note: not all stakeholders are shown.

Figure 6-3 Eight GDP stakeholder groups

Stakeholder Group	Definition
Air Transportation Service Providers (ASPs)	The collection of commercial, competitive, for-profit enterprises that operate and provide scheduled passenger and cargo air transportation service to airports in the U.S.
Airports	The operators of airports used for commercial flight service and that may provide other essential and non-essential services for airlines and their passengers.
Business	The at-large business community that benefits directly or indirectly from air transportation.
FAA	The Federal Aviation Administration: an administration of the Department of Transportation, under the Executive branch of the U.S. Federal Government; responsible for the oversight of the aviation industry in the U.S. Includes the traffic manager, who is responsible for designing a GDP and the Terminal Radar Approach Control (TRACON) facility that manages the arrivals at an airport.
Federal Government	The Executive branch of the Federal Government of the U.S., other than the FAA; appoints and oversees the work of the FAA.
Local/Regional/State Governments	Local, regional, and State governments that may be responsible for the oversight of airports or benefit from the commercial impact of aviation
Passengers	Connecting and terminating passengers on scheduled flights that may use capacity-constrained resources of the National Airspace System (NAS)
Society	Society-at-large

These eight groups will be used as the basis for the stakeholder analysis presented in this chapter. The analysis is conducted at a high level of organization for two reasons. First, it is more likely that an analysis performed at a high level will represent all groups that are relevant to the design of a GDP. For example, GDPs will invariably impact a myriad of smaller air transportation companies, such as air taxis and fractional jets, but an explicit enumeration of these small carriers would require a larger analysis. Second, many of the stakeholders, especially smaller ones, may have a relatively low salience. By performing the analysis at a higher level, it is possible to determine which groups of stakeholders will have higher salience. The results of this analysis could be used to target research efforts to explore groups of stakeholders that are more relevant to a GDP.

Section 6.1.2 Stakeholder Salience

The salience of each stakeholder (or, in this case, stakeholder group) is determined using a framework described by Mitchell et al. (1997). This framework classifies stakeholders based upon three attributes. The first is power, which refers in this case to the ability of a stakeholder to cause changes to the design of a GDP. The second, legitimacy, is the degree to which a GDP affects a stakeholder. The third attribute is urgency, which measures the need perceived by a stakeholder to change the GDP. The difference between urgency and legitimacy is perception: a stakeholder that is urgent but not legitimate perceives a need to change a GDP regardless of whether there is an actual need; and a stakeholder that is legitimate but not urgent is affected by a GDP but may not be aware of it. The more attributes exhibited, the more salient the stakeholder; the most salient stakeholder is one with the need (legitimacy), willingness (urgency), and ability (power) to influence the design of a GDP.

The Mitchell classification is both appropriate and useful for this analysis because 1) it is assumes that stakeholders are defined broadly and 2) may be applied to qualitative data. Consistent with Mitchell's description of the analysis, each of the three attributes is assumed to be binary and each stakeholder is assumed to possess each attribute to either a "high" or "low" degree. Many of the stakeholders identified in the previous section will only possess one or two of these attributes to a high degree. A feature of the Mitchell framework is that it will classify a stakeholder as one of eight possible types, as based upon the combination of attributes that it possesses to a high degree (Figure 6-4). For brevity, a discussion of the various types will be postponed to the results of the analysis in §6.4.

Figure 6-4 Stakeholder classifications (Mitchell (1997))

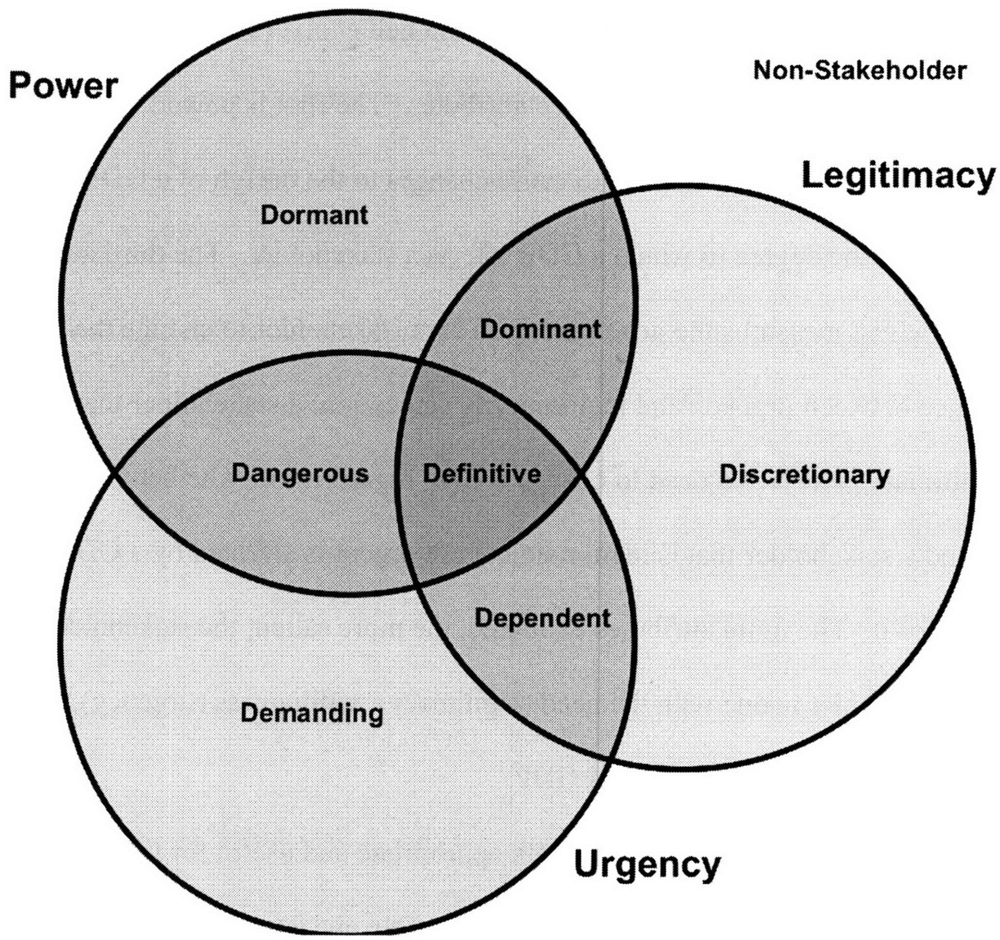

In the context of a GDP, power, legitimacy, and urgency are defined as follows:

- Power: The ability to influence the FAA and/or the outcome of a GDP
- Urgency: The degree to which a stakeholder perceives the need to change the GDP
- Legitimacy: The degree to which a GDP may affect a stakeholder

The power, legitimacy, and urgency of each stakeholder are determined by two analyses: the first examines the transactions between the various stakeholders and the second determines the potential impact of a GDP upon each stakeholder. Stakeholder power, which is defined by the first analysis, is discussed in §6.2.2; urgency and legitimacy are discussed after the second analysis in §6.3.4 and §6.3.5, respectively.

Section 6.1.3 Influencing the Outcome of a GDP

The previous discussion introduces the idea that a stakeholder may be able, in some way, to affect a GDP. Before continuing, it is helpful to consider both the types of changes that are possible in a GDP and the mechanism by which change is produced. GDPs are introduced in §1.1 as tools used by traffic managers to manage the arrival demand at an airport by assigning ground delay to aircraft. However, in practice, the specific decisions made by the traffic manager are only part of a larger process.

The *GDP Design Process* is defined as the set of three steps. In the first step, a traffic manager gathers information and opinions in regards to forecasts of the arrival demand and capacity of an airport from meteorologists, local air traffic control centers, and commercial airlines. This solicitation is formalized through a series of conference calls held every morning between these groups. In the second step, the traffic manager enacts a GDP, assigning an initial set of ground delays to aircraft. In the third step, these initial delays may be revised to reflect flight cancellations and popup flights (additional arrival demand). During the third step, commercial airlines are also permitted to adjust the delays assigned to their flights through two formalized mechanisms: slot-credit-substitutions (SCS) and flight swaps[42].

Figure 6-5 Three steps of the GDP design process

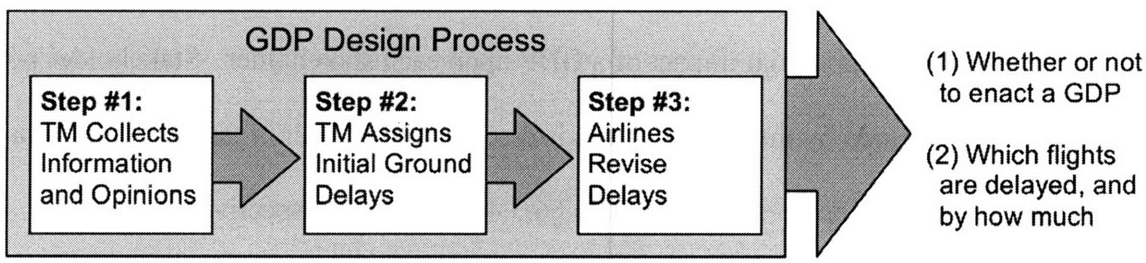

There are two means by which a group may influence a GDP. The first is by acting through the GDP design process, either by providing information before the initial assignment of ground delay in step 1 or by adjusting the assigned delays after step 3. This type of change to the process is limited to the airlines (ASPs) and the traffic manager (FAA), which are the two stakeholders that have formalized control of the decision process. The second means of influencing the outcome of a GDP is to change the process, itself.

The adaptation of the GDP design process highlights the value of broadly defining the notion of a stakeholder. The current body of research in GDPs has largely focused on how to design a GDP subject to changes within the process. Instead, by working at a broader level, it is also possible to consider changes to the process, involving more stakeholders and opening a new area in which to explore the design and use of GDPs.

[42] For more detail, please refer to Volpe (2004).

Section 6.2 Stakeholder Transactions

The transactions between the various stakeholders are important for two reasons. First, the power of each stakeholder is a measure of the ability of a stakeholder to influence the FAA or the outcome of a GDP. This influence is a result of something the group provides to the FAA, such as money or information. Second, transactions are also the means by which the costs of delays – and the impacts of a GDP – are transmitted to each of the stakeholder groups. Thus, the network of transactions also determines, in part, stakeholder legitimacy.

Transactions between the stakeholder groups may exist at several levels. For example, the FAA is allocated money from the Federal Government for its operational budget. Although part of this allocation may be used to fund, in various manners, the design of a GDP, it applies to all functions of the FAA. In contrast, airlines interact directly with the traffic manager to provide flight schedule information for a new GDP or to revise the delays assigned by an existing program. Thus, while the analysis presented in this chapter focuses on the relationships that exist at the organizational level, it also includes interactions at a more detailed level that are specific to the design of a GDP.

Section 6.2.1 A General Description of Stakeholder Relationships

A detailed accounting of the transactions between the various stakeholders is presented in the appendix. Across these relationships, there are three general processes that drive the transactions.

The first process is the flow of information to and from the FAA during the GDP design process. Although airlines provide specific data to the FAA, such as flight schedules, airline representatives may also provide advice and suggestions regarding the possible design or modification of a program directly to the traffic manager. Airlines also communicate with passengers regarding the flights that are delayed by a GDP. For their part, passengers provide the airline with information both during a GDP, such as willingness to take an alternate itinerary, and after, through feedback given to customer service representatives. Figure 6-06 shows a map of the key information flows between stakeholders, with those specific to the design of a GDP highlighted in bold.

The second key process is the flow of goods and services related to the air transportation industry. In particular, airlines and their passengers engage in economic activities with other stakeholders. For example, airlines purchase various goods and services from other businesses and employ pilots, cabin crew, and ground personnel. Considered together, the U.S. commercial air carriers directly employ more than half million full time workers domestically[43]. Passengers also participate in additional economic activities, either directly though the purchase of rental cars, hotel rooms, etc., or indirectly through business transactions that are facilitated by travel.

Airports play a unique role in the flow of goods and services because of their co-dependence on other groups related to the provision of air transportation. Airports provide essential infrastructure in the form of runways and terminals for the air transportation industry. Other goods and services that are required by the industry, such as fuel, food service, cleaning, etc., may also be provided by an airport operator or by a third party through a concession. At a fundamental level, airports facilitate the transactions between airlines, their passengers, and other businesses that serve the needs of both.

[43] 2006 employment for certified carriers, U.S. Bureau of Transportation Statistics

Figure 6-6 Flow of information between stakeholders

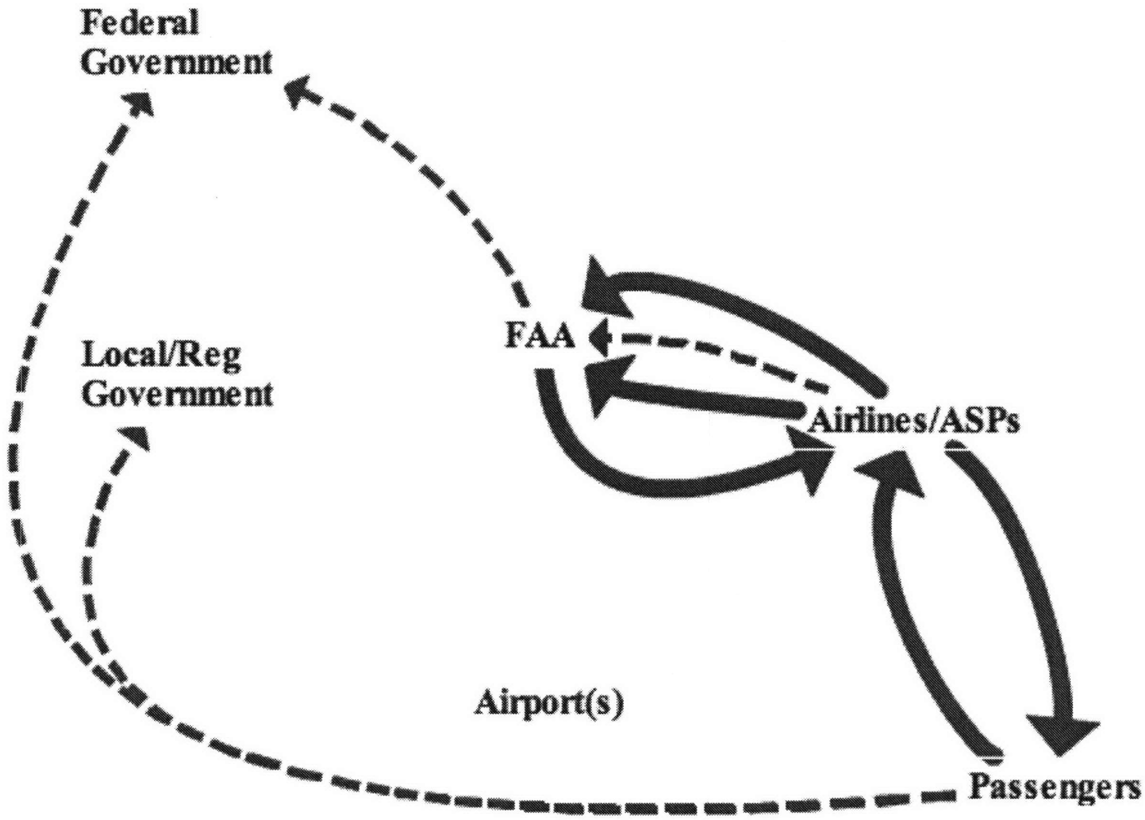

In exchange, airport operators receive payment for the use of their facilities by both airlines (e.g. rent) and passengers (e.g. PFCs), and may also profit from other economic activity by selling fuel to airlines and concessions for passenger and employee amenities (food courts, stores, parking, etc.). Airports also provide access to air transportation for local residents and businesses; these positive externalities may be a reason why many municipalities and regional governments invest in the construction of airports.

The role of the FAA in this system is to provide a management service to ASPs and airports, ensuring the safe and efficient use of the NAS. From the perspective of the FAA, the economic benefits of air transportation are realized in the form of operational funding. The FAA uses two principal funding sources: the first is the Federal

Government, which collects corporate and personal income taxes and distributes some to the FAA, and the second the Aviation Trust Fund (ATF), which is funded by a tax on each commercial aviation ticket. Although airlines and passengers provide funding to the FAA through the ATF, their contributions are required by law, as mandated by the Federal Government.

Figure 6-7 and Figure 6-8 show the general flow of financial capital and goods and services between the various stakeholders in regards to air transportation. Of particular note, passengers are both an indirect source of the financial capital that reaches the FAA and an indirect end recipient of the services provided by the FAA. However all interactions between passengers and the FAA are through intermediary stakeholders, such as the ASPs.

The third process is the legal control that one stakeholder may have over another. This control may be in the form of political power or represented as an enforceable legal contract. For example, both the Federal Government and the State, Local, and Regional Governments are elected by their constituents, which form the general society. In turn, the Federal Government appoints the Administrator of the FAA, who is ultimately responsible for the FAA's policies. Although society does not control the day-to-day operations or decisions made by the Federal Government or the FAA, society may act to replace the government if there is disapproval of the decisions that are made (or lack thereof). As an example, in November 2007, President Bush spoke of the need to reduce air transportation delays[44], an action which represents both a directive to the FAA, as well as a message to society that the Federal Government is working to reduce delays.

[44] "President Bush discusses aviation congestion," White House press release, November 15, 2007

Figure 6-7 Flow of financial capital between stakeholders

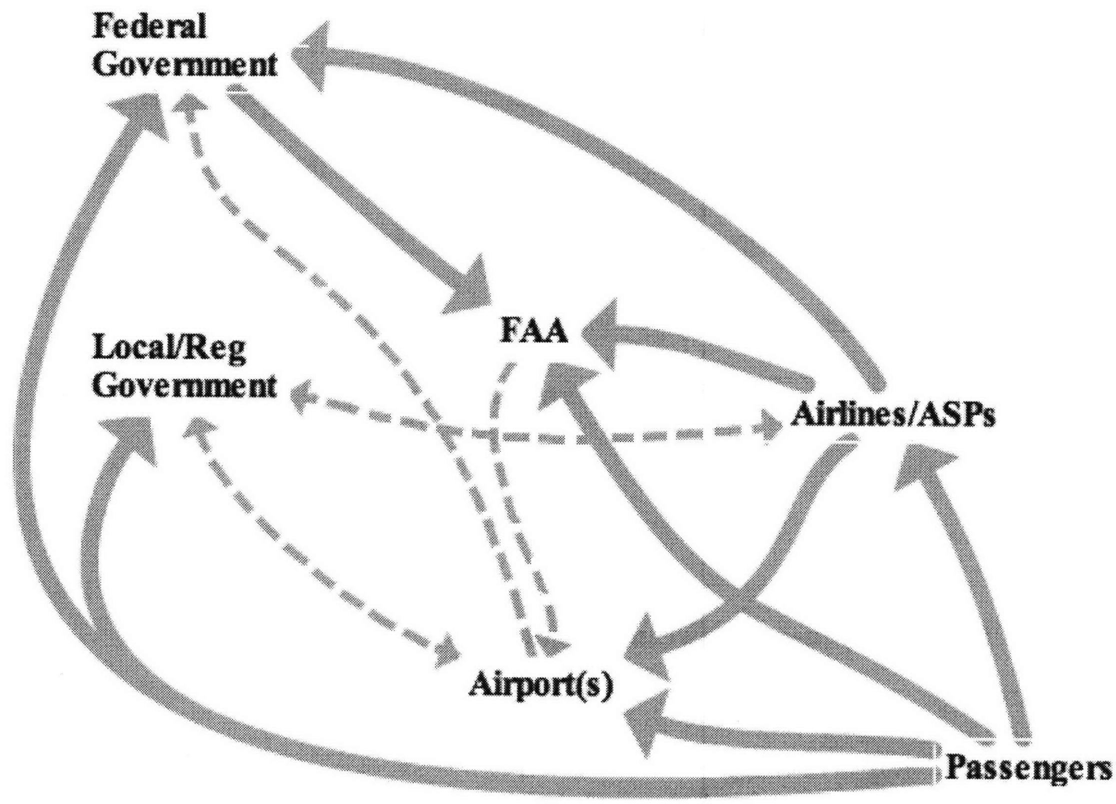

Figure 6-8 Flow of goods and services between stakeholders

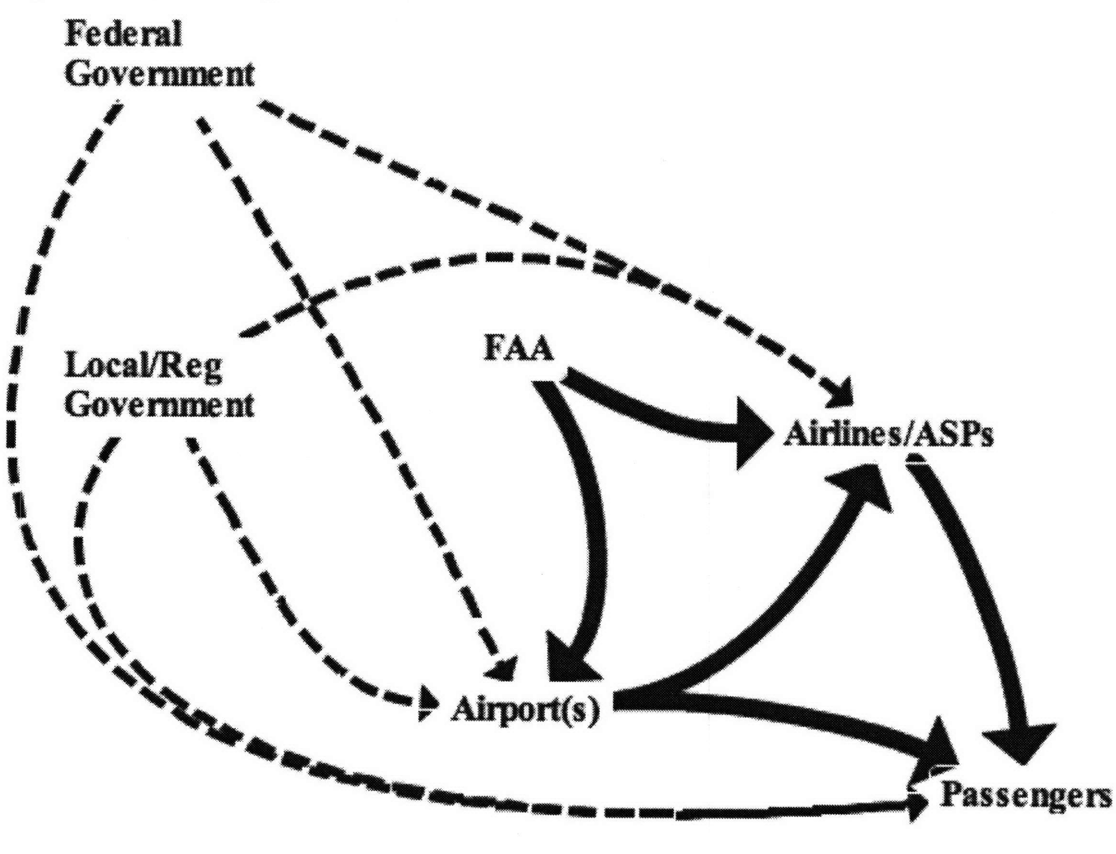

The FAA exerts a form of control over both airports and ASPs by providing for regulation, certification, and oversight of the industry, participating businesses, and its infrastructure. A very recent example of the power of the FAA is the cancellation of more that 3,000 American Airlines flights and delays to 250,000 passengers between April 8 and 12, 2008, due to FAA inspections of wiring harnesses on MD-80 aircraft[45]. More specific to the design of a GDP, the FAA assigns controlled departure times to aircraft as part of a GDP.

Figure 6-9 Flow of mandates between stakeholders

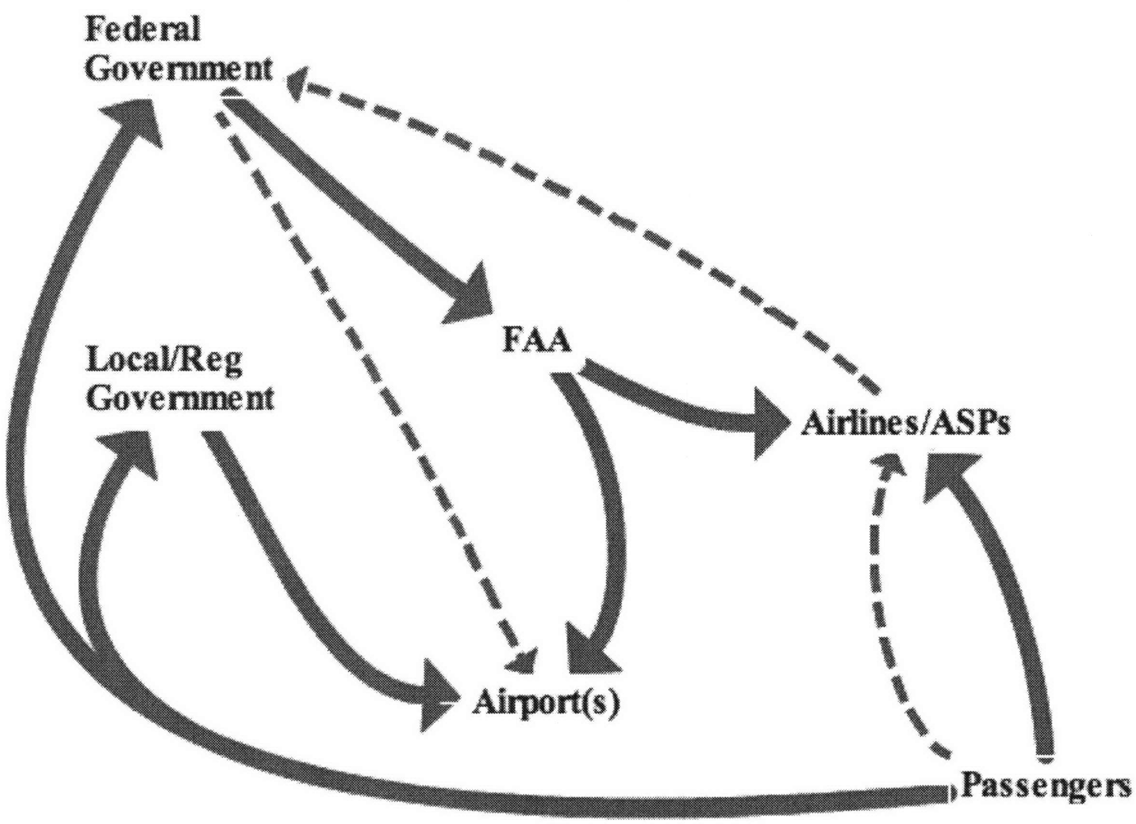

[45] "Bush keeping 'a close eye' on airline passenger misery, holds cabinet meeting Monday" International Herald Tribune, April 11, 2008

Passengers also hold a form of legal control over the airlines as represented by the contract formed by the purchase of a ticket. Airlines are obligated to provide transportation once this contract is formed (subject to various stipulations). Even passengers who are bumped, or denied boarding, from their scheduled flight are entitled to compensation by law[46]. Figure 6-9 shows a flow of legal obligations, or mandates, through the various stakeholders.

Section 6.2.2 Stakeholder Power

The power of a stakeholder is represented by the various relationships that the stakeholder has with the other groups. It is through these relationships that the stakeholder gains information, acts upon another stakeholder in order to enforce its will, or is acted upon. For the design of a GDP, stakeholder groups that interact directly with the FAA will, generally, have higher power than those that are only tangentially related. Specifically, recall that power is defined as "the ability to influence the FAA and/or the outcome of a GDP." The FAA meets this definition because the FAA can control the outcome of a GDP. Other stakeholders with high power are those that provide a necessary input to the FAA, such as money or information, or that exert some measure of legal control over the FAA.

The two other stakeholders with high power are the Federal Government and the ASPs (Figure 6-10). First, the Federal Government has high power because it has a measure of political control over the FAA and because it is directly responsible for funding the operations of the FAA. Second, ASPs have high power because of the

[46] New Federal rules announced April 16, 2008 raise the maximum compensation to $800 for domestic passengers.

information they provide to the traffic manager and their direct involvement in the GDP design process. Other stakeholders, such as passengers and airports have low power because they do not exert influence directly over the FAA or the design of a GDP.

Figure 6-10 Stakeholder power

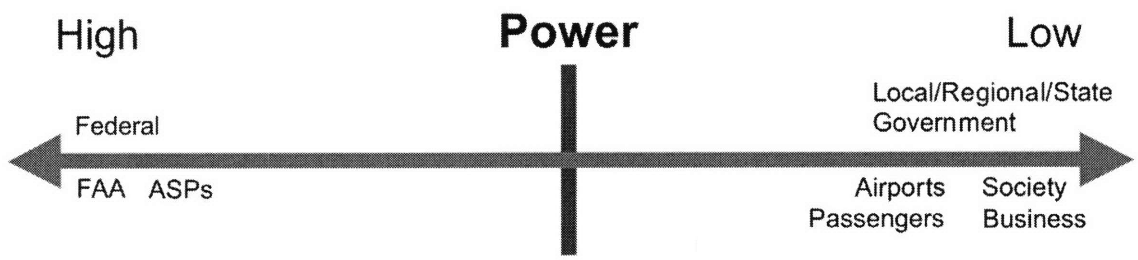

To facilitate the discussion of stakeholder legitimacy and urgency, it is necessary to first discuss the impacts of a GDP, both actual and perceived; legitimacy is discussed in §6.3.4 and urgency in §6.3.5.

Section 6.3 Impacts of a GDP

At a basic level, a GDP assigns ground delay to aircraft in order to avoid or reduce the possibility of future airborne delays. However, the impact that a GDP will have depends not only on the stakeholder, but also on the perspective from which the delays are observed. This section discusses the impacts of a GDP from three perspectives and the impacts that each has on the various stakeholders

Section 6.3.1 A GDP from the Flight Perspective

A GDP assigns ground delay to particular aircraft in order to reduce the airborne delays of other aircraft. On the whole, the affect is assumed to be positive, because ground delay is preferable to airborne delay for reasons of cost and efficiency. However, from the perspective of individual flights, the distribution of assigned delay is important – some flights will benefit from a GDP while others will be negatively impacted. The decisions made during the GDP design process will affect the nature of this distribution. For a discussion of how this distribution may be affected, please refer to §3.1.2 and Vossen (2002).

The distribution of delay is especially important from the perspectives of ASPs and of passengers. Commercial airlines (and other ASPs) are competitive entities that provide transportation and a level of service for customers. Each ASP operates a group of flights; if there is a correlation between the flights operated by an ASP[47] and the relative benefit offered to a flight by a GDP, it might give a carrier a competitive advantage. Similarly, for passengers, the distribution of delay will affect which passengers gain or lose as a result of a GDP.

For other stakeholders, the distribution of delay has a much smaller impact. For example, from the perspective of the airport to which the GDP is applied, flights will be delayed before arrival regardless of whether delay is taken on the ground or in the air. However, the effect is not negligible because a GDP can only increase the overall amount of delay to the flights at the airport. Furthermore, as airborne delay results in more fuel, which many airports sell to the airlines, a GDP may even have a negative effect. However, airports, as well as the other stakeholders, are related economically to the ASPs; thus, the impact of a GDP on these stakeholders is most likely to be tied into how the GDP treats the airlines.

[47] Such as that identified by Hanowsky (2007)

Section 6.3.2 A GDP from the ASP Perspective

The assignment of delay, itself, that is caused by a GDP has a second, more subtle effect on the operations of larger commercial airlines. Under nominal operating conditions, without a GDP, an airline is free to adjust its operational schedule at any time for any reason. For example, the departure of a flight may be delayed to accommodate the loading of additional baggage and passengers or to allow for repairs to non-essential parts, such as the entertainment system or a lavatory. However, flights that are delayed on the ground by a GDP must depart within a given window of time or risk being assigned additional delay. Thus, a second impact of a GDP is a reduction in the flexibility that airlines have over their own schedules.

It is assumed that this loss of control and flexibility reduces the efficiency of the operations of an airline. However, the extent to which it impacts an airline, and whether it impacts other ASPs, has not been studied in the GDP literature. Therefore, for the purposes of this analysis, it is assumed that ASPs and their passengers feel the primary impact, with other stakeholders possibly receiving a small economic effect.

Section 6.3.3 A GDP from the Air Transportation Network Perspective

A third effect of a GDP is the reduction of airborne congestion throughout the air transportation network. Air transportation networks are complex, large-scale, and highly-connected. A key emergent property of these networks is that the congestion caused by airborne delays in one part of the network may quickly spread to other areas. For example, severe weather over Tennessee might result in delays for flights from Boston to Chicago, even though these flights would not pass through or near the weather. Although a GDP will not reduce the arrival delay time for the airport to which it is applied, it can serve a critical role in reducing airborne congestion and the spread of delays through the network. Furthermore, when viewed strategically, congestion poses a significant threat to the air transportation industry. Severe chronic delays, or network gridlock could change the competitive landscape or render reliable service untenable.

Reducing airborne congestion has a significant and positive impact on various stakeholders. For ASPs, passengers, and airports there is a reduction in airborne delay. For the FAA, reducing airborne congestion reduces the workload of air traffic controllers. Other stakeholders benefit through their relationships with passengers, airports, and ASPs. And, strategically, all stakeholders benefit from a functioning air transportation system.

Section 6.3.4 Legitimacy

Recall that the legitimacy of a stakeholder is a measure of how the decisions made by the GDP design process impacts that stakeholder. As highlighted by the previous discussion, it is clear that all stakeholders, to some degree, are impacted by GDPs. Taken to an extreme, without the use of any GDPs, the resulting severe delays and congestion might decrease the supply of air transportation in the U.S. However, the level of impact that GDPs have varies by stakeholder group.

The groups that are most directly impacted are ASPs and their passengers. Not only are aircraft and passenger delays a direct result of a GDP, but these stakeholders also benefit greatly from widespread reductions in airspace congestion. The cost of a delay may vary by aircraft, such that ASPs and passengers may also be very sensitive to how delays are distributed by a GDP. Furthermore, airlines may also be sensitive to how a

GDP may impact the manner by which they manage their operations. The reduction in schedule flexibility that results from the initiation of a GDP may result in operational inefficiency for airlines and wider economic impacts for other stakeholders.

The FAA is also directly impacted by a GDP. As a method of demand management, GDPs may reduce air traffic congestion and the peak workload for an air traffic controller (as measured in number of managed aircraft). Peak workload is important because there may be safety concerns if the number of aircraft managed by a controller exceeds operational thresholds. Other stakeholders are also affected by the design of a GDP, but to a lesser extent. For each of these, the primary concern is that air traffic congestion reduces the overall efficiency of the NAS. Severe congestion may lead to wider economic inefficiencies, which first affect business and society, and consequently, all levels of government.

In order to apply the Mitchell framework, it is necessary to divide the stakeholders into two groups, those with high legitimacy (such as ASPs, passengers, and the FAA) and those with low legitimacy (Federal Government and Local, Regional, and State Governments). Three stakeholder groups (airports, society and business) are affected by a GDP, but only to a limited extent, and require a more detailed discussion.

For airports, which generate revenue based upon the business activities of airlines and passengers, the economic impacts are mixed. To a certain extent, the choice of which flights are delays may impact the passenger throughput (and revenues) of the airport. However, a GDP will not reduce the overall arrival delay across all flights, thus the direct impact on an airport may be small. The overall effects of GDPs on the air transportation industry are likely to have greater impact on the revenues of airports. As a result, the legitimacy of airports depends on the overall efficiency of the NAS; if the efficiency would be greatly reduced without GDPs, then the air transportation industry would suffer, and the impact on airports would be significant.

A similar argument may be made for society and business; that if the economic livelihood of these two groups would be adversely affected if GDPs were not used, then both have high legitimacy. The author is not aware of any studies that have sought to explore how the NAS might perform without the use of GDPs; however, given the current levels of congestion and delays, it is assumed that the congestion would be significantly worse without GDPs. Thus, airports, business, and society are assumed to have high legitimacy.

Figure 6-11 Stakeholder legitimacy

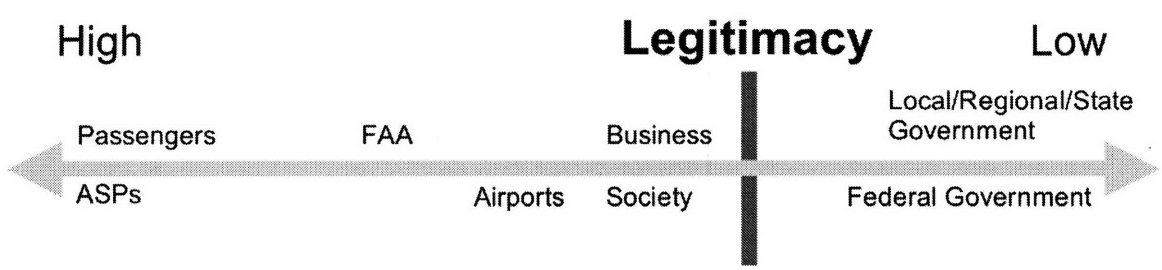

Section 6.3.5 Stakeholder Urgency

As mentioned previously, urgency measures the need that is perceived by the stakeholder to change the GDP. Although the urgency of a stakeholder may increase with the magnitude of the impact of a GDP on that stakeholder, urgency is a measure of perception; stakeholders that are unaware of the impact of a GDP upon them, regardless of magnitude, will have low urgency, and those that perceive a need to act will have high urgency. Consequently, those stakeholders that are more directly related to a GDP, such as those that deal with GDPs on a daily basis, will be more aware of GDPs and will also have a higher urgency.

The stakeholders with the highest urgency are those that perceive a direct effect from a GDP. These stakeholders are the ASPs and passengers, which are assigned delays directly by a GDP, and the FAA, which acts to initiate a GDP. An airport is also an urgent stakeholder to the extent that a GDP affects the flights and passengers at that airport.

For other stakeholders, urgency may be in flux. While air transportation delays have been perceived as a problem by the air transportation industry for many years, the Federal Government has only recently pushed for a national solution to delays. The recent comments made by President Bush (§6.2.1) and attention of the media to delays at airport, suggests that the perception of the impacts of congestion and delays – and, as a result, GDPs – may be increasing. Thus, while these other stakeholders are shown in Figure 6-12 as having low urgency, this status may change.

Figure 6-12 Stakeholder urgency

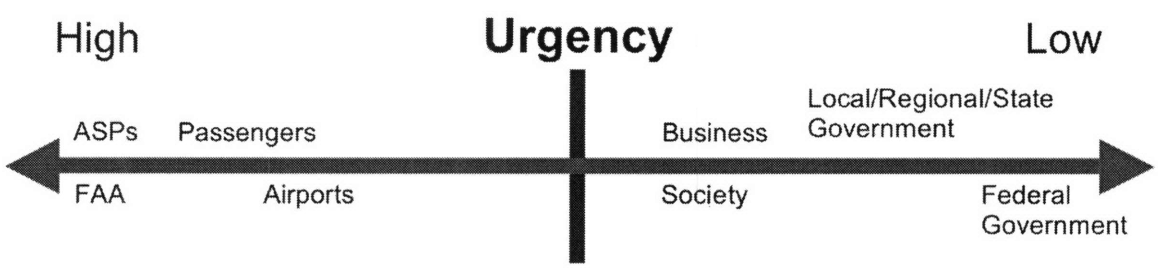

Section 6.4 Stakeholder Salience

Given the measures of power, legitimacy, and urgency, as defined in the preceding sections, the application of the Mitchell framework is straight-forward. The salience of each stakeholder is determined by the degree to which it possesses the power, legitimacy, and urgency attributes. Recall that the Mitchell classification requires that each attribute be classified on a binary scale for each stakeholder; for this analysis, each stakeholder-attribute is described as either "high" or "low" (Figure 6-13).

For the design of a GDP, the most salient stakeholders are the Air Transportation Service Providers and the FAA, both of which are affected by the outcome of a GDP and have the ability and the willingness to influence the process. This result follows what is observed in practice; the traffic manager and various air carriers are able to and do participate in the decisions to initiate a GDP.

Figure 6-13 Stakeholder classifications

Stakeholder	Power	Urgency	Legitimacy	Classification
Air Transportation Service Providers	High	High	High	Definitive
Airports	Low	High	High*	Dependent
Business	Low	Low	High*	Discretionary
FAA	High	High	High	Definitive
Federal Government	High	Low	Low	Dormant
Local/Regional/State Governments	Low	Low	Low	Non-Stakeholder
Passengers	Low	High	High	Dependent
Society	Low	Low	High*	Discretionary

Note: * assumes that the use of GDPs result in a significant reduction to NAS congestion

Figure 6-14 Stakeholder classification presented as a Venn diagram

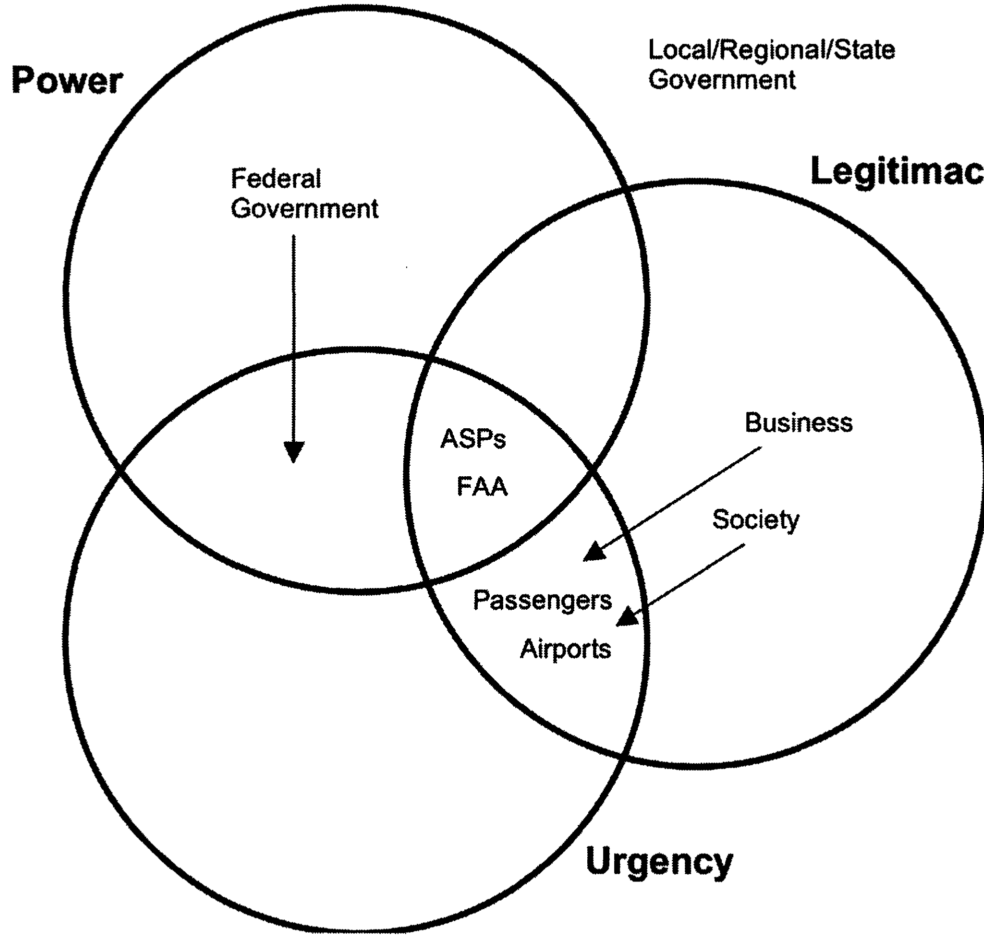

Each of the other stakeholders exhibits high levels for some, but not all three, of the attributes, which reduces the impetus for them to be considered during the design process. For example, passengers and airports are affected by GDPs and might make changes to their design, but have no power to do so. Business and Society are also affected by GDPs but do not perceive an urgent need to act. On the other hand, the Federal Government has the power to influence the FAA, but may also not perceive the need to act. Lastly, Local/Regional/State Governments exhibit low values for each of the attributes and are only tangentially related to the design of a GDP.

The classifications provided by the analysis may be applied to determine which of the stakeholders should be considered during the design of a GDP. For example, both ASPs and the FAA have a high value for all three attributes, indicating a high degree of salience. Therefore, the design of a GDP not only considers the needs of the FAA and ASPs, but it should continue to do so. However, other stakeholders may also deserve consideration as based on how they might become more salient. For example, the discussion of stakeholder urgency highlights the potential for several stakeholders, including the Federal Government, business, and society to become more urgent. This potential is indicated by the small arrows in Figure 6-14.

The Federal Government is of particular interest because it is also a powerful stakeholder. The Mitchell classification describes powerful and urgent stakeholders as "dangerous;" the danger exists in the sense that the Federal Government might use its power to direct the FAA to reduce perceived problems that might otherwise have little impact on the effects of a GDP. For example, the Federal Government opened up military airspace off the east cost of the U.S. in order to reduce congestion during the busy Thanksgiving holiday. However, this measure did little to relieve congestion around the busy New York airports or at the major passenger hubs in Atlanta, GA (ATL); Chicago, IL (ORD); and Dallas, TX (DFW).

Further analysis also suggests that the needs of passengers might deserve greater consideration during the design of a GDP because they could become empowered. The analysis presented in §6.2 highlights the relationship of passengers with other stakeholders. In particular, passengers exert some form of control over both the Federal Government and ASPs; the first in the form of an election process (via society) and the second as a result of the contract formed by the purchase of a ticket. Thus, although passengers do not have power directly over the FAA, they may have power over stakeholders that do. In contrast, airports (which are also described as a "dependent" stakeholder of the FAA), would be less likely to gain the support of powerful stakeholders and are, therefore, less salient.

Section 6.5 Conclusions

The objective of the analysis presented in this chapter is to identify the stakeholders that should receive consideration during the design of a GDP. Nine stakeholder groups are identified and evaluated in regards to three attributes: power, legitimacy, and urgency. The analysis shows that three stakeholders, the FAA, ASPs, and passengers, should receive more consideration than other stakeholders during the design of a GDP. This section presents two general conclusions from the analysis.

Section 6.5.1 Stakeholder Consideration in Practice

The first conclusion is that the design of GDPs in practice may not consider the appropriate objectives. Of the three salient stakeholder identified by the analysis, two are currently involved in the GDP design process (Figure 6-5, page 158) used in practice. The third stakeholder group, passengers, does not directly participate and, therefore, may not be adequately considered.[48]

Furthermore, the analysis of GDP impacts also identifies a potential conflict between the interests of ASPs and those of other stakeholders. The initiation of a GDP

[48] The analysis in Chapter Seven will demonstrate that the current process is significantly less than optimal from the passengers' perspective.

reduces airlines' operational flexibility, which is assumed to have a cost to each ASP. However, GDPs reduce airborne congestion, which is beneficial for all stakeholders. Thus, there may exist situations in which an ASP, which has the ability to influence the GDP design process, would oppose a GDP that would be beneficial to the stakeholders as a whole.

Section 6.5.2 Areas for Future Research

A second conclusion drawn from the analysis is that further research could be conducted to explore the salience of the organizations that comprise each stakeholder group. For example, ASPs is comprised of the legacy passenger carriers, "low cost" and point-to-point carriers, air taxis, general aviation, etc. (Figure 6-2, page 153). The power, legitimacy, and urgency of each of these groups may differ from that of the group as a whole, so that a stakeholder analysis could identify which airlines are important. Other groups for which a deeper analysis could help are passengers and airports.

Additionally, stakeholder theory could also be applied to understand the concept of equity in regards to a GDP. Equity is a term that has begun to appear in the academic literature[49] in reference to the distribution of ground delay. Specifically defined, however, an equitable allocation is "consonant with merit or importance"[50] and therefore requires an *a priori* understanding of what each aircraft, or airline, etc. deserves. This is the essence of stakeholder salience.

Finally, additional research could be conducted to explore the ways in which passengers might be considered in the design of a GDP. The discussion of stakeholder

---

[49] §2.2
[50] Merriam-Webster online dictionary, April 23, 2007

salience (§6.4) suggests that both airlines and the Federal Government could empower passengers; the airlines by using their influence as part of the existing GDP Design Process and the Federal Government by changing the nature of the process, itself. The next chapter presents a comparison of two methods that represent these alternatives, to test which method results in a more favorable outcome for the passenger and their implications for other stakeholders.

Chapter 7 Ground Delay Programs and the Passenger

The stakeholder analysis discussed in Chapter Six highlights various impacts that flight delays may have on different stakeholder groups and the role of GDPs in alleviating these impacts. One group that is particularly affected is the passengers, who may be delayed for extended periods of time with little or no control over the delays that affect them. In regards to the assignment of ground delay, passengers have little or no power because they are neither involved in the design process nor do they have any direct influence over the FAA. However, the needs of passengers could become more relevant to the design of a GDP if other, more powerful stakeholders, such as the airlines or the Federal Government, were to act on their behalf.

This chapter presents and compares two alternative approaches through which passengers could be considered during the design of a GDP. The first is based on current practice in which airlines are permitted to adjust the ground delays that are initially assigned by the traffic manager. This approach assumes that each airline will use an existing mechanism called a flight swap to minimize the total delay cost experienced by its own passengers. The second method hypothesizes that the GDP design process could be revised so that the traffic manager is responsible for assigning delays that minimize the total passenger delay cost. The key difference between these two approaches is the level at which decisions considering the passenger are made; in the former, the airline considers the welfare of the passenger, while in the latter, passengers of all airlines are considered by the traffic manager.

Chapter Seven is presented in four sections. First, §7.1 defines passenger delay cost in regards to the design of a GDP. Second, §7.2 uses an assignment model to determine how airlines might be able to reduce the cost of delays to their passengers through the use of flight swaps. Third, §7.3 proposes a new MIP formulation for the design of a stochastic and dynamic GDP in which flight delay costs may vary by aircraft. This new model is applied to the SAGHP to represent a GDP in which the traffic manager minimizes the total passenger delay cost. Finally, §7.4 compares the results of these two approaches for a hypothetical problem. This comparison suggests that the total passenger delay cost of a GDP might be significantly reduced if the traffic manager were to consider passengers during the initial design of a GDP.

Section 7.1 A Model of Passenger Delay Cost

Before proposing methods of reducing the cost of delays to passengers, it is necessary to define passenger delay cost in regards to a GDP. A model of delay cost introduced in §3.1 proposes that there are three types of delay, ground, airborne, and cumulative, and that flights incur cost as a function of the duration of each type. The *flight delay cost* of each flight is defined as the sum of these costs; the total flight delay cost of a GDP is the sum of the flight delay costs across all flights. For convenience and to be consistent with this model, the *passenger delay cost* incurred by a flight is defined as the flight delay cost multiplied by the number of passengers[51] on the flight. Therefore,

[51] For the examples discussed in this chapter, the number of seats on each aircraft is used as a proxy for the number of passengers

the total passenger delay cost of a GDP is the sum of the passenger delay costs that are incurred by each flight.

Thus, if a flight f incurs costs GDC_f, ADC_f, and CDC_f as a result of ground, airborne, and cumulative delay, respectively, then the flight delay cost $Cost_f$ is given as:

$$Cost_f = GDC_f + ADC_f + CDC_f \tag{7.01}$$

And, if f carries pax_f passengers, then the passenger delay cost $PaxCost_f$ is given by:

$$PaxCost_f = pax_f \times Cost_f \tag{7.02}$$

In order to differentiate the flight delay costs from the passenger delay costs of a GDP, flight delay cost is reported in "ac-units" and passenger delay cost in "pax-units."

This definition makes two important assumptions regarding the manner in which passengers incur delay cost. First, it assumes that all of the passengers on each flight accrue cost in an equal manner. In reality, each passenger may have a unique function to express the cost of a given delay. For example, a business traveler may be very sensitive to short delays while a leisure traveler might be less sensitive. This simplification is made to facilitate the analysis presented in this chapter and could be relaxed in future work; all of the models proposed in this chapter could be adapted to consider cost functions that vary by passenger.

The second assumption is that the cost experienced by passengers on a given flight is independent of the delays assigned to other flights. In reality, some of the passengers that are delayed as part of a GDP may be connecting to another flight, in which case the cost due to a delay in the arrival of the first flight might be negligible if the passengers make the connection and significant if the passengers do not. However, including passenger connections is non-trivial from both a practical and a modeling perspective. During the course of a GDP, airlines may adjust their flight schedules and reroute aircraft, so it may be difficult to determine the length of delay that would cause a passenger to miss a connection. Consideration of connecting passengers is suggested as an area of future research.

On the other hand, the assumption that passenger delay costs are independent of flight connections is convenient for several reasons. First, it is not only consistent with the design of GDPs in practice, in which the traffic manager does not explicitly consider the connections of aircraft and crew, but it could also be estimated using data that is currently available to the traffic manager. Second, this definition is consistent with existing models in the stochastic/dynamic SAGHP literature, such as the ECM. These models are based on the premise that the total delay cost of a program is the sum of the costs incurred by individual flights. By maintaining this consistency, these models may be more easily adapted to solve for ground delays that minimize the total passenger delay cost.

Section 7.2 Passenger Interest as Represented by the Airlines

The objective of this analysis is to test the hypothesis that the passenger delay cost of a GDP could be reduced if the initial assignment of delays made by the traffic manager were to consider the number of passengers on each flight. The test, itself, will compare the total expected passenger delay cost for two hypothetical GDPs. Both use the same forecasts of arrival capacity and demand, but differ in the manner in which passengers are considered during the design process.

The first GDP, or "base case approach," assumes that the traffic manager assigns delays in order to reduce the total expected flight delay cost for a stochastic and dynamic forecast of the airport arrival capacity and that airlines may adjust these delays using flight swaps in order to minimize the total passenger delay cost. The flight swap mechanism and its limitations will be discussed in greater detail in the next section. The second GDP, or "alternative approach," assumes that the traffic manager assigns delays that minimize the total expected passenger delay cost. Figure 7-1 summarizes the two approaches.

Figure 7-1 Two approaches for reducing the passenger delay cost of a GDP

	Base Case Approach	**Alternative Approach**
Traffic Manager	Assigns delay that minimizes the total expected flight delay cost	Assigns delay that minimizes the total expected passenger delay cost
Airlines	Revise the delays to minimize the passenger delay cost for each airline	Take no action to revise ground delays.

Section 7.2.1 The Flight Swap Mechanism

In the base case approach, ground delays are assigned in two steps. First, the traffic manager assigns a controlled demand time to each flight. This represents the time at which the flight will land if there is sufficient arrival capacity at the airport. Consistent with a GDP, the difference between this controlled time and the *scheduled* arrival time of the flight is assigned as ground delay[52]. The overall assignment of flight arrival demand times is assumed to minimize the total flight delay cost, as is the case in the solution to the ECM.

[52] §3.2.2

In the second step, each airline revises the ground delays assigned by the traffic manager using a mechanism called a flight swap. A flight swap permits an airline to exchange the arrival demand times, or slots, between two of the aircraft that it operates. It is assumed in this section that each airline will swap flights in order to minimize the total delay cost of its own passengers. (In practice, an airline will likely use flight swaps to minimize its own delay costs, which might not minimize the cost to passengers.)

A key limitation of this approach is that flight swaps are limited in how they may reassign arrival slots. First, the use of flight swaps might not increase the cumulative airport arrival demand for any period of time. Second, flight swaps may only be used in conjunction with flights from the same airline. And, third, flight swaps cannot assign a flight to depart earlier than *originally* scheduled. Thus, while flight swaps may reduce the passenger delay cost for both individual airlines and the GDP, as a whole, the use of swaps will likely not minimize the total passenger delay cost across all flights.

Figure 7-2 shows an example of a flight swap. Consider five flights that are scheduled to arrive at an airport that is subject to a GDP (column #1). Each aircraft is operated by one of two airlines, 01A and 02A. In the first step, the traffic manager assigns a controlled arrival time to each flight (column #2). Note that flights #003 and #005 are assigned to times 1504 Z and 1508 Z, resulting in two and four minutes of ground delay, respectively. In step #2, airline 01A swaps the times assigned to flight #003 and #005 (column #3), which reduces the ground delay of #005 to zero minutes and increases the ground delay of #003 to six minutes. Note that this is the only feasible swap given the flights in this example.

Figure 7-2 An example of a flight swap

Time	(#1) Scheduled Flight Arrivals	(#2) Flight Arrivals Assigned by GDP	(#3) Flight Arrivals After Swap
1500 Z	01A-001	01A-001	01A-001
1501 Z	02A-002		
1502 Z	01A-003	02A-002	02A-002
1503 Z	02A-004		
1504 Z	01A-005	01A-003	**01A-005**
1505 Z			
1506 Z		02A-004	02A-004
1507 Z			
1508 Z		01A-005	**01A-003**

In order to determine the reduction in passenger delay cost that could be achieved using flight swaps, a new optimization model is proposed. This section presents a linear programming (LP) formulation that solves for an allocation of flight arrival demand times to aircraft that minimizes the total cost of passenger delays for each airline. The model, which may be used in a manner consistent with the current use of fight swaps, is referred to as the "Airline Assignment Model," or AAM.

Section 7.2.2 The Airline Assignment Model LP Formulation

The Airline Assignment Model (AAM) is a linear programming formulation that identifies an assignment of aircraft to arrival slots that minimizes the total assignment cost. Unlike the ECM, which is used to create a GDP, the AAM assumes that a GDP has already been created and that it is optimal in terms of the total flight delay cost. The AAM may be used to consider factors that have not been included in the design of the original GDP. For example, if the ECM designs a GDP that minimizes the total expected aircraft delay cost, the AAM could be applied to the solution to the ECM in order to minimize the total passenger delay cost.

Before defining the formulation, it is helpful to define "slot" in a manner that is consistent with a stochastic and dynamic arrival capacity forecast. Recall that stochastic and dynamic forecasts are assumed to be provided as a set of discrete profiles, each of which indicates the airport arrival capacity over time and that the profiles are grouped into a scenario tree[53]. Let there exist a GDP for a set of flights I and an arrival capacity scenario tree with a set of profiles Q. The arrival demand times assigned to each flight as part of the GDP may vary by profile. So, for each flight $i \in I$, there exists a slot j that consists of the set of the arrival demand times assigned to i across Q. Let the set of all slots be J, such that $|I| = |J|$ by definition.

Although each flight is assigned to a slot as part of the initial GDP, it might be possible to assign each flight to one of many slots. It is assumed that the set of possible assignments, which may be limited by practical considerations, and the cost of each possible assignment are known *a priori*. Let the feasibility of an assignment be given by F_{ij}, which equals 1 if the assignment of flight i to slot j is feasible and 0 if the assignment is infeasible, and the cost of assigning flight i to slot j by C_{ij}. As defined, F_{ij} and C_{ij} allow the AAM to make a generic assignment of flights to arrival slots; both will be discussed further in regard to flight swaps and passenger delay costs in §7.2.2.

The decision variable for the AAM is a set of binary variables $X_{ij} \in \{0,1\}$ that indicates the assignment of flights to slots.

[53] §2.3.1

$$X_{ij} = \begin{cases} 1 & \textit{if flight i is assigned to slot j} \\ 0 & \textit{otherwise} \end{cases}$$

The objective function and constraints of the AAM are defined as follows:

Minimize: $$\sum_{I,J} C_{ij} \times X_{ij} \qquad (7.03)$$

Subject To: $$\sum_{J} X_{ij} = 1 \qquad \forall \quad i \in I \qquad (7.04)$$

$$\sum_{I} X_{ij} = 1 \qquad \forall \quad j \in J \qquad (7.05)$$

$$X_{ij} \leq F_{ij} \qquad \forall \quad i \in I; \; j \in J \qquad (7.06)$$

The objective function (7.03) minimizes the total cost of the assignment of flights to slots. Constraints (7.04) and (7.05) ensure that each flight is assigned to exactly one slot and that exactly one flight is assigned to each slot. Constraints (7.06) specify that only feasible assignments may be made. Provided that there exists at least one feasible assignment (which is given by the initial GDP), the set of constraints may be represented as a totally unimodular matrix and the AAM may be solved as an LP relaxation.

Section 7.2.3 Adaptation of the AAM to Minimize Passenger Delay Costs

As formulated, the AAM makes a generic assignment of flights to arrival slots. However, by specifying values for F_{ij} and C_{ij}, the AAM can be used to identify the set of flight swaps that would minimize the total expected passenger delay cost.

Feasible Flight Assignments

For the assignment of a flight i to a slot j to be feasible, it must be consistent with the use of a flight swap, which is defined as meeting four conditions:

1. The airline that operates flight i must be the same as the airline that operates the flight assigned to slot j by the initial GDP
2. Each arrival demand time that corresponds to slot j must be equal to or greater than the *scheduled* arrival time of i
3. The departure times for i that would correspond to each arrival demand time for j must not violate the coupling, or non-anticipativity, constraints (3.15) of the ECM
4. If i is exempt from the GDP, then the only slot to which it may be assigned is that from the initial GDP

The third condition is especially important and unique to a swap that would occur for a slot defined for a stochastic and dynamic arrival capacity forecast. As determined by the optimal solution to the ECM, the departure times of a flight under various profiles are conditional on the times at which the arrival capacity forecast is revised[54]. For example, if a forecast of the arrival capacity indicates that (1) the future capacity will be either high or low and (2) a revised forecast at 1500 Z will indicate which of the two will occur; then it is not possible for a flight to depart at 1430 Z under the high-capacity profile and 1445 Z under the low-capacity profile.

[54] §3.2.4

The Cost of a Flight Assignment

As defined in §7.1, passenger delay cost is the number of passengers on a flight times the flight delay cost, which is a function of the airborne, ground, and cumulative delay time that would be experienced by that flight. For a GDP subject to uncertain arrival capacity, the objective is to minimize the total *expected* passenger delay cost. Therefore, let

$$C_{ij} = \sum_{Q}\left(p_q \times PaxCost_{ijq}\right) \qquad \forall \quad i \in I;\ j \in J \tag{7.07}$$

Where $PaxCost_{ijq}$ represents the passenger delay that results from assigning flight i to arrival slot j under arrival capacity profile q and p_q is the likelihood of q (§3.2.1). The value of $PaxCost_{ijq}$ can be derived from the number of passengers and scheduled arrival time of flight i, the scheduled landing times of slot j, and a set of flight delay cost functions, such as those in §3.3.1, all of which are assumed to be known *a priori*.

Section 7.2.4 Passenger Delay Costs for the Base Case GDP

As a demonstration, the base case approach is applied to a simple, hypothetical example. For brevity and convenience, the initial GDP that is required by the AAM is taken from the solution to the ECM for a previous example from §3.3.3. The existing solution to the ECM is convenient because it not only provides arrival demand and landing times for each slot, but it also minimizes the total expected flight delay cost. As an additional input for this example, aircraft carrier and type data are taken from the ETMS database.

A sample of ETMS data is shown in Figure 3-7 (page 60). The carrier of each flight is given by the first three characters of the aircraft ID (ACID)[55]. Aircraft type (TYPE) is used to estimate the number of passengers on each flight. For this analysis, aircraft are assumed to be 100% full, so the number of passengers on each flight is equal to the number of seats in a typical configuration for the given aircraft type. For example, a typical configuration of an Airbus 319 (A319) has 124 seats[56]. It is assumed that airlines would have accurate passenger count data in practice.

Figure 7-3 shows the total number of aircraft and passengers by carrier. For these data, there are 348 flights and 31 carriers. However, 262 of the scheduled flights, or 75%, are operated by only four carriers. This suggests that many of the carriers will have little or no opportunity to swap flights. Furthermore, opportunities to reduce passenger delays using swaps will be limited for two of the larger carriers (10A and 28A) because they operate fleets with aircraft of uniform, or nearly uniform, size[57].

Figure 7-3 also shows a summary of the aircraft and passenger delay costs by airline, for the initial GDP as defined by the solution to the ECM. Specifically:

1. Total expected flight delay cost (in ac-units)
2. Total expected flight delay time (in ac-periods)
3. Total expected passenger delay cost (in pax-units)
4. Total expected passenger delay time (in pax-periods)

The total expected passenger delay cost of this GDP is 36,785.7 pax-units and the total expected passenger delay time is 31,584.8 pax-periods[58].

[55] The true carrier identities are available in ETMS data but have been masked here by three-letter codes.
[56] A complete table of passenger counts by aircraft type is provided in the appendix.
[57] Recall that for this hypothetical example, all aircraft are assumed to be 100% full.
[58] Delay time is measured as the total cumulative delay time.

Figure 7-3 Total aircraft and passenger costs for the ECM

Airline	Count of Flights	Count of Pax	Avg. Pax per Flight	Total Exp. Flight Delay Cost	Total Exp. Flight Delay Time	Total Exp. Pax Delay Cost	Total Exp. Pax Delay Delay
01A	68	11,119	163.5	67.0	56.4	10,481.3	8,817.6
03A	3	298	99.3	0.2	0.1	24.8	12.4
08A	5	290	58.0	1.2	1.2	70.0	70.0
09A	1	150	150.0	1.0	1.0	150.0	150.0
10A	50	2,500	50.0	42.5	38.2	2,125.0	1,910.0
11A	2	493	246.5	0.0	0.0	0.0	0.0
12A	1	305	305.0	0.0	0.0	0.0	0.0
13A	1	253	253.0	0.0	0.0	0.0	0.0
14A	2	87	43.5	1.3	1.3	61.1	61.1
16A	1	452	452.0	1.3	1.3	587.6	587.6
17A	1	100	100.0	0.0	0.0	0.0	0.0
18A	12	824	68.7	6.5	6.5	434.2	434.2
20A	4	583	145.8	6.0	4.2	916.8	627.9
21A	2	100	50.0	0.7	0.7	35.0	35.0
24A	3	463	154.3	3.7	3.7	571.5	571.5
26A	2	648	324.0	0.2	0.1	24.8	12.4
27A	1	130	130.0	2.2	1.3	286.0	169.0
28A	69	3,587	52.0	42.5	41.5	2,252.1	2,182.1
34A	1	50	50.0	2.2	2.1	110.0	105.0
38A	1	305	305.0	1.2	0.3	366.0	91.5
39A	1	524	524.0	0.0	0.0	0.0	0.0
40A	7	350	50.0	4.8	3.0	240.0	150.0
43A	1	50	50.0	1.3	1.3	65.0	65.0
44A	3	407	135.7	0.0	0.0	0.0	0.0
46A	2	310	155.0	1.5	1.5	232.5	232.5
47A	3	355	118.3	9.7	9.7	1,128.5	1,128.5
48A	2	976	488.0	6.5	6.5	3,305.2	3,305.2
49A	1	335	335.0	0.0	0.0	0.0	0.0
50A	18	1,000	55.6	15.9	15.9	873.0	873.0
53A	75	12,378	165.0	71.9	57.9	11,882.3	9,451.1
54A	5	605	121.0	5.2	5.0	563.0	542.2
TOTAL	**348**	**40,027**	**115.0**	**296.5**	**260.7**	**36,785.7**	**31,584.8**

Results

Figure 7-4 (page 192) shows a summary of the delays that are assigned by the solution to the AAM. Applying the assignment model reduces the total passenger delay cost by 3.5% and the total passenger delay time by 4.3% across all airlines as compared to the solution to the ECM. Furthermore, there is no change to the overall number of flight arrivals over time as specified by the GDP. Thus, flight swaps reduce the passenger delay cost while maintaining the same amount of aircraft delay time.

However, the results also show that most individual carriers do not benefit from the opportunity to use flight swaps. Only three carriers (01A, 28A, and 53A) are able to use swaps to reduce their passenger delay costs. Furthermore, of these, only 53A achieves a reduction in cost that is greater than 5%[59].

Figure 7-4 Aircraft and passenger costs for the ECM + AAM

	Count of		Total Exp. Flight Delay		Total Exp. Pax Delay			
Airline	Flights	Pax	Cost	Time	Cost	Chg.	Delay	Chg.
01A	68	11,119	67.0	56.4	10,314.8	-1.6%	8,650.5	-1.9%
03A	3	298	0.2	0.1	24.8	0.0%	12.4	0.0%
08A	5	290	1.2	1.2	70.0	0.0%	70.0	0.0%
09A	1	150	1.0	1.0	150.0	0.0%	150.0	0.0%
10A	50	2,500	42.5	38.2	2,125.0	0.0%	1,910.0	0.0%
11A	2	493	0.0	0.0	0.0	0.0%	0.0	0.0%
12A	1	305	0.0	0.0	0.0	0.0%	0.0	0.0%
13A	1	253	0.0	0.0	0.0	0.0%	0.0	0.0%
14A	2	87	1.3	1.3	61.1	0.0%	61.1	0.0%
16A	1	452	1.3	1.3	587.6	0.0%	587.6	0.0%
17A	1	100	0.0	0.0	0.0	0.0%	0.0	0.0%
18A	12	824	6.5	6.5	434.2	0.0%	434.2	0.0%
20A	4	583	6.0	4.2	916.8	0.0%	627.9	0.0%
21A	2	100	0.7	0.7	35.0	0.0%	35.0	0.0%
24A	3	463	3.7	3.7	571.5	0.0%	571.5	0.0%
26A	2	648	0.2	0.1	24.8	0.0%	12.4	0.0%
27A	1	130	2.2	1.3	286.0	0.0%	169.0	0.0%
28A	69	3,587	42.5	41.5	2,142.2	-4.9%	2,072.2	-5.0%
34A	1	50	2.2	2.1	110.0	0.0%	105.0	0.0%
38A	1	305	1.2	0.3	366.0	0.0%	91.5	0.0%
39A	1	524	0.0	0.0	0.0	0.0%	0.0	0.0%
40A	7	350	4.8	3.0	240.0	0.0%	150.0	0.0%
43A	1	50	1.3	1.3	65.0	0.0%	65.0	0.0%
44A	3	407	0.0	0.0	0.0	0.0%	0.0	0.0%
46A	2	310	1.5	1.5	232.5	0.0%	232.5	0.0%
47A	3	355	9.7	9.7	1,128.5	0.0%	1,128.5	0.0%
48A	2	976	6.5	6.5	3,305.2	0.0%	3,305.2	0.0%
49A	1	335	0.0	0.0	0.0	0.0%	0.0	0.0%
50A	18	1,000	15.9	15.9	873.0	0.0%	873.0	0.0%
53A	75	12,378	71.9	57.9	10,855.6	-8.6%	8,361.4	-11.5%
54A	5	605	5.2	5.0	563.0	0.0%	542.2	0.0%
TOTAL	**348**	**40,027**	**296.5**	**260.7**	**35,482.6**	**-3.5%**	**30,218.1**	**-4.3%**

[59] It should be noted that the results of the AAM represent the lower bound of passenger delays cost that is achievable to an airline using flight swaps; if an airline selects a different set of swaps, then the reduction in passenger delay cost will likely be less than that found by the AAM.

Section 7.3 Passenger Interests as Represented by the Traffic Manager

The second, or alternative, approach is proposed as a contrast to the current process and demonstrates the potential for reducing passenger delay costs if the traffic manager were to design a GDP that minimizes passenger delay costs. As compared to the airlines, which use flight swaps to minimize the delay costs to their own passengers, the traffic manager could tailor the initial design and assignment of the arrival slots to minimize the total passenger delay cost across all flights.

In order to consider passengers during the design of a GDP, it is necessary to model delay costs as varying by flight. Unfortunately, the ECM is limited to considering delay cost functions that are the same for all aircraft. Therefore, this section proposes a new MIP formulation, which solves for the optimal assignment of ground delay given cost functions that are non-linear and may also vary by aircraft. This new formulation is applied to the same example used to demonstrate the AAM in the previous section. A comparison of the results shows that the total expected passenger delay cost could be significantly reduced if the traffic manager were to consider passengers during the design of the initial program.

Section 7.3.1 The Flight-Defined Cost MIP Formulation

The model presented in this section solves the SAGHP subject to a stochastic and dynamic arrival capacity forecast, a capacity of the airborne arrival queue, and non-linear delay cost functions that may vary by aircraft. This formulation is very similar to the ECM except that it enhanced in order to consider cost functions that may vary by flight; for this reason, it is referred to as the Flight-Defined Cost Model, or "FDCM."

The decision variables for the formulation of the FDCM are the same as those previously defined for the ECM in §3.2.2.

$$d_{ftq} \in \{0,1\} \qquad \forall \; f \in F; \; t \in \tau; \; q \in Q$$

$$\delta_{ftq} \in [0,1] \qquad \forall \; f \in F; \; t \in \tau; \; q \in Q$$

$$\lambda_{ftq} \in [0,1] \qquad \forall \; f \in F; \; t \in \tau; \; q \in Q$$

The variables d_{ftq}, δ_{ftq}, and λ_{ftq} represent the fraction of flight f that departs, demands arrival or lands, respectively, by time t under profile q.

The objective function and constraints (7.08) – (7.23) for the FDCM are also the same as those in the ECM, with two key exceptions. First, the costs of cumulative, ground, and airborne delay in constraints (7.09) – (7.13) are expressed as functions of incremental cost parameters that are specific to each aircraft. It is assumed that these parameters $\left(\chi_t^f, \gamma_t^f, \alpha_t^f\right)$ are provided as an input to the formulation.

The second difference is the addition of flight arrival coupling constraints (7.23), which state that the flight arrival indicator variables for each flight must be the same under all profiles in each scenario. For the FDCM, in which airborne delay cost may vary by aircraft, it is possible for the optimal order of flight arrivals to depend upon the future airport arrival capacity. Coupling constraints ensure that the decisions regarding which flights land in each period may only consider the information that would be available during the time period when the decision is made.

As a special case, if all aircraft have an identical, non-decreasing airborne delay cost function, then the objective function is minimized by a first-come-first-serve arrival order, regardless of which profile is ultimately realized. This occurs because, among the flights in the arrival queue, the flight that has been in queue longest also has the greatest marginal cost of airborne delay. However, the use of flight-specific cost functions in the FDCM means that a first-come-first-serve order may no longer minimize the total expected delay cost. Deviation from this order in the solution to the FDCM will be discussed further in §7.3.2.

FDCM Objective Function and Constraints

$$\textit{Minimize} \qquad \sum_{Q}\left[p_q \times \sum_{F}\left(CDC_{fq} + ADC_{fq} + GDC_{fq}\right)\right] \tag{7.08}$$

$$CDC_{fq} = \sum_{t=SAT_f}^{T-1}\left(\left(1-\lambda_{ftq}\right)\times \chi^{f}_{t-SAT_f+1}\right) \qquad \forall\ f \in F;\ q \in Q \tag{7.09}$$

$$GDC_{fq} = \sum_{t=SDT_f}^{T+SDT_f-SAT_f-1}\left(\left(1-d_{ftq}\right)\times \gamma^{f}_{t-SDT_f+1}\right) \qquad \forall\ f \in F;\ q \in Q \tag{7.10}$$

$$ADC_{fq} = \sum_{t=SAT_f}^{T-1}\left(\left(1-\lambda_{ftq}\right)\times \alpha^{f}_{t-SAT_f+1}\right) - \sum_{s=0}^{T-SAT_f-1}\ \sum_{t=s}^{T-SAT_f-1} Z_{fstq} \qquad \forall\ f \in F;\ q \in Q \tag{7.11}$$

$$Z_{fstq} \le \left(\alpha^{f}_{t-s+1} - \alpha^{f}_{t-s}\right)\times\left(1-\lambda_{f(t+SAT_f)q}\right) \tag{7.12}$$

$$Z_{fstq} \le \left(\alpha^{f}_{t-s+1} - \alpha^{f}_{t-s}\right)\times\left(1-d_{f(s+SDT_f)q}\right) \tag{7.13}$$

$$\forall\ f \in F;\ s \in 0\ldots T-SAT_f-1;\ t \in s\ldots T-SAT_f-1;\ q \in Q$$

$$d_{ftq} \ge d_{f(t-1)q} \qquad \forall\ f \in F;\ t \in 1\ldots T;\ q \in Q \tag{7.14}$$

$$\lambda_{ftq} \ge \lambda_{f(t-1)q} \qquad \forall\ f \in F;\ t \in 1\ldots T;\ q \in Q \tag{7.15}$$

$$d_{f(t-STE_f)q} = \delta_{ftq} \qquad \forall\ f \in F;\ t \in STE_f\ldots T;\ q \in Q \tag{7.16}$$

$$\lambda_{ftq} \le \delta_{ftq} \qquad \forall\ f \in F;\ t \in \tau;\ q \in Q \tag{7.17}$$

$$\delta_{ftq} = 0 \qquad \forall\ f \in F;\ t \in 1 \ldots SAT_f - 1;\ q \in Q \tag{7.18}$$

$$\lambda_{fTq} = 1 \qquad \forall\ f \in F;\ q \in Q \tag{7.19}$$

$$\sum_F \left(\delta_{ftq} - \lambda_{ftq}\right) \le W_{tq}^{MAX} \qquad \forall\ t \in \tau;\ q \in Q \tag{7.20}$$

$$\sum_F \left(\lambda_{ftq} - \lambda_{f(t-1)q}\right) \le M_{tq} \qquad \forall\ t \in \tau;\ q \in Q \tag{7.21}$$

$$d_{ftq_i} = \ldots = d_{ftq_j} \qquad \forall\ f \in F;\ b \in B;\ q_i, \ldots, q_j \in b;\ t \in \{b_s \ldots b_e\} \tag{7.22}$$

$$\lambda_{ftq_i} = \ldots = \lambda_{ftq_j} \qquad \forall\ f \in F;\ b \in B;\ q_i, \ldots, q_j \in b;\ t \in \{b_s \ldots b_e\} \tag{7.23}$$

Section 7.3.2 An Application of the FDCM

The FDCM is applied to the sample problem discussed in §7.2.3. For consistency, the same experimental setup is maintained; except that the incremental delay cost parameters now reflect the passenger delay cost. Let:

$$\chi_t^f = pax_f \times \chi_t \tag{7.24}$$

$$\gamma_t^f = pax_f \times \gamma_t \tag{7.25}$$

$$\alpha_t^f = pax_f \times \alpha_t \tag{7.26}$$

Note that this change is consistent with the definition of passenger delay cost in §7.1.

Results

The total expected cost of the solution to the FDCM is 24,002.6 pax-units and the total expected passenger delay time is 17,840.6 pax-periods. These values represent a 34.8% reduction in the total passenger delay cost and a 43.5% reduction in the total passenger delay time as compared to the solution to the ECM (Figure 7-5, page 198).

These percentage reductions are also an order of magnitude greater than those achieved by the combined ECM/AAM solution.

This significant reduction in total passenger delay cost may require an increase in total flight delays and delay cost. The solution to the FDCM model is 1.1% higher in total expected flight delay cost and 1.4% higher in total expected flight delay time (measured as cumulative flight delay) as compared to the original ECM solution. Thus, a GDP that is optimal in terms of passenger delay cost may assign more flight delays than one that optimizes flight delay cost. However, for this example, the increase in flight delay cost is small in comparison to the reduction in passenger delay cost. If the relative difference shown for this example were to prove typical of what would occur in practice, then the slight increase in flight delay cost might be viewed as acceptable to the airlines and to the traffic managers.

A second issue is that the solution to the FDCM favors those airlines that operate larger aircraft and may result in smaller cost reductions or cost increases for aircraft with fewer passengers. For example, the FDCM increases the total expected passenger delay cost of airline 28A by 77% over the solution to the ECM. Clearly, such an impact raises the issue of equity for the carriers and passengers that would be disadvantaged by this approach[60]. The analysis presented in §6.4 identifies airlines as a *definitive* stakeholder group, with the power to influence the FAA and to control the outcome of a GDP. Therefore, future research should examine the salience of individual airlines and the relative impact of equity.

---

[60] Especially as many of the disadvantaged passengers may be connecting to other flights

Figure 7-5 Results of the FDCM

Airline	Count of Flights	Pax	Total Exp. Flight Delay Cost	Time	Total Exp. Pax Delay Cost	Chg.	Delay	Chg.
01A	68	11,119	26.6	13.6	4,099.6	-61%	2,090.9	-76%
03A	3	298	0.2	0.1	24.8	0%	12.4	0%
08A	5	290	1.0	1.0	50.0	-29%	50.0	-29%
09A	1	150	0.0	0.0	0.0	-100%	0.0	-100%
10A	50	2,500	84.6	84.2	4,230.0	+99%	4,210.0	+120%
11A	2	493	0.0	0.0	0.0	0%	0.0	0%
12A	1	305	0.0	0.0	0.0	0%	0.0	0%
13A	1	253	0.0	0.0	0.0	0%	0.0	0%
14A	2	87	4.0	4.0	192.2	+215%	192.2	+215%
16A	1	452	0.0	0.0	0.0	-100%	0.0	-100%
17A	1	100	0.0	0.0	0.0	0%	0.0	0%
18A	12	824	5.0	4.9	272.4	-37%	264.6	-39%
20A	4	583	4.4	2.6	706.2	-23%	417.3	-34%
21A	2	100	1.2	1.2	60.0	+71%	60.0	+71%
24A	3	463	0.2	0.1	31.0	-95%	15.5	-97%
26A	2	648	0.2	0.1	24.8	0%	12.4	0%
27A	1	130	2.2	1.3	286.0	0%	169.0	0%
28A	69	3,587	83.4	83.2	3,976.9	+77%	3,962.9	+82%
34A	1	50	4.4	4.4	220.0	+100%	220.0	+110%
38A	1	305	0.6	0.2	183.0	-50%	61.0	-33%
39A	1	524	0.0	0.0	0.0	0%	0.0	0%
40A	7	350	17.3	17.3	865.0	+260%	865.0	+477%
43A	1	50	1.5	1.5	75.0	+15%	75.0	+15%
44A	3	407	0.0	0.0	0.0	0%	0.0	0%
46A	2	310	0.0	0.0	0.0	-100%	0.0	-100%
47A	3	355	7.5	7.4	862.5	-24%	851.0	-25%
48A	2	976	2.4	1.4	1,243.2	-62%	726.4	-78%
49A	1	335	0.0	0.0	0.0	0%	0.0	0%
50A	18	1,000	18.0	17.8	944.0	+8%	930.0	+7%
53A	75	12,378	31.7	14.8	5,276.4	-56%	2,296.2	-76%
54A	5	605	3.4	3.2	379.6	-33%	358.8	-34%
TOTAL	**348**	**40,027**	**299.8**	**264.3**	**24,002.6**	**-34.8%**	**17,840.6**	**-43.5%**

Finally, although these results show that the FDCM may have the potential to significantly reduce the passenger delay cost of a GDP, this has only been demonstrated for a single, hypothetical example. Further experimentation with the FDCM should be an area of future research. However, there are two reasons to believe that the substantial reduction achieved by the FDCM for this example is indicative of reductions that may be encountered in practice. First, the flight arrival demand for this example is representative of the actual demand for a metropolitan airport (§3.3.1). Second, the strategy of giving preference to larger aircraft is used in practice on a smaller scale; airlines often postpone or cancel flights with fewer passengers in order to reduce delays to larger aircraft.

Section 7.4 Conclusions

The experiment presented in this chapter tests the hypothesis that the passenger delay cost of a GDP could be reduced if passengers were to be considered by the traffic manager during the initial assignment of ground delay. Two different approaches are presented for reducing the total expected passenger delay cost. The first uses a two-step process that initially assigns ground delays in order to minimize the total expected flight delay cost and then uses flight swaps to minimize the total passenger delay cost of each airline. This approach is based on current practice, in which the airlines are assumed to consider the needs of passengers, and uses two models, the ECM and AAM, to identify the optimal assignment of delay in each step.

The second approach is an alternative to current practice and assumes that the traffic manager would consider passengers directly during the design of a program. This approach minimizes the total expected passenger delay cost in a single step with a new model, the FDCM. When applied to a simple, hypothetical example, the second, or alternate, approach achieves a much greater reduction in passenger delay cost than is possible given the first approach (Figure 7-6, page 200). However, the second approach may increase the delay costs to some airlines and would require changes to the policies underlying the current process by which GDPs are designed and implemented in practice. This section examines the practicality of a passenger-optimal GDP and suggests areas for future research.

Figure 7-6 A comparison of two approaches that reduce passenger delay cost.

	ECM	ECM + AAM	% Chg	FDCM	% Chg
Flight Cost (ac-units)	296.5	296.5	+ 0.0%	299.8	+ 1.1%
Flight Delay (ac-periods)	260.7	260.7	+ 0.0%	264.3	+ 1.4%
Pax Cost (pax-units)	36,785.7	35,482.6	– 3.5%	24,002.6	– 34.8%
Pax Delay (pax-periods)	31,584.8	30,218.1	– 4.3%	17,840.6	– 43.5%

Section 7.4.1 Implications of the FDCM for the Design of GDPs in Practice

The experimental results suggest that if passenger delay cost is considered during the initial design of a GDP by the traffic manager, then there is the potential for a significant reduction in the total expected passenger delay cost. However, the alternative approach requires several key assumptions regarding the design of GDPs that might preclude its use in practice.

First, by considering passengers, a traffic manager would give preference to airlines that fly larger aircraft. For the example shown earlier, this preference offers significant benefits to airlines 01A and 53A, both of which achieve a greater than 75% reduction in passenger delay cost, at the expense of those that fly smaller aircraft, such as 10A and 28A. This preference is contrary to the current design of GDPs, which assumes that all of the non-exempted aircraft should be treated equally.

Existing work in the SAGHP literature on equity by Vossen (2002), Vossen and Ball (2003), and Mukherjee (2004), is consistent with this assumption and measures the "fairness" of a GDP with reference to the amount of ground delay that is assigned to flights. However, defining equity only in terms of assigned ground delay does not capture many elements of the problem, some of which were brought to light by the stakeholder analysis in Chapter Six. Therefore, it is strongly suggested that further research on the equity of GDPs be conducted. A possible starting point might be to extend the stakeholder analysis in the previous chapter to determine the salience of individual airlines and the relative cost or benefit to these airlines of considering passenger costs during the design of a GDP. Additional work might also consider equity in regards to other stakeholder groups, such as passenger market segments, airports, local communities, etc.

Second, as previously noted, the solution to the FDCM may deviate from the first-come-first-serve (FCFS) order assumed by prior SAGHP models. As part of the optimal solution, the FDCM specifies the order in which flights should arrive: flights with more passengers may be permitted to jump ahead of smaller aircraft in the arrival queue. Although flight arrivals, in practice, may not always occur in FCFS order, the policies of air traffic control may preclude the preferential treatment suggested by the solution. For the example discussed in this chapter, the optimal solution to the FDCM assigns an FCFS arrival order to all but four aircraft. However, as the model is only demonstrated for a single example, it is not known whether this is the norm or an exception. Therefore, further research should be conducted to test the effects of the non-FCFS assumption and ascertain its practicality.

In conclusion, this chapter suggests that significant reductions in passenger delay cost could be achieved if the traffic manager were to consider passengers during the initial assignment of flight delays. However, in order to achieve these reductions in practice, the benefits offered to other powerful stakeholders, such as airlines, would need to offset the costs. Further research is necessary to identify how the changes suggested by the FDCM would impact different stakeholders and the salience of each stakeholder in the design of a GDP. As a final thought, given the billions of dollars spent annually on increasing the capacity of the NAS, if a large reduction in passenger delays is possible with the existing capacity of the system, there is reason to believe that such change could be achieved. At the very least, even if the barriers to change are too difficult to overcome, the FDCM could be used to inform the FAA and other stakeholders of the hidden cost of these barriers.

Section 7.4.2 Thoughts on Directing the Development of Future SAGHP Models

The role of Chapter Seven is two-fold. Inasmuch as the analysis presented here demonstrates how the cost of delays to passengers could be considered during the design of a GDP, this chapter also ties the stakeholder analysis presented in Chapter Six to the more quantitative concepts discussed earlier in the thesis. The FDCM is built upon the ECM, itself an extension of earlier SAGHP models. However, while the advances made by each successive generation of models has been toward reflecting more accurately the realities of GDPs as currently practiced – from uncertainty, to dynamics, to models of cost – the FDCM is, intentionally, less tied to existing assumptions about GDPs.

Specifically, the purpose of the FDCM is to test a fundamental assumption that is inherent to the current design of GDPs in practice. By allowing delay costs to vary by aircraft, passenger delay costs can be considered directly by the traffic manager during the design of a program. The motivation for this test is provided by an analysis of the stakeholders of a GDP, which identifies passengers as a dependent stakeholder. Moving forward, the results of the FDCM suggest that further research should be conducted to examine the salience of individual airlines and classes of passengers. And, as a result of this work, it may be possible to further tailor how passengers are considered in order to benefit passengers while not affecting significantly the benefits that other stakeholders derive from the use of GDPs.

Chapter 8 A Summary and Conclusions

This dissertation proposes a mixed-integer programming model for solving the single-airport ground holding problem (SAGHP) subject to a stochastic and dynamic forecast of the airport arrival capacity. This model, the ECM, offers two features not found in previous stochastic and dynamic SAGHP models: 1) a wider range of delay cost functions and 2) a capacity of the airborne arrival queue. As compared to previous models from the literature, the solution to the ECM either results in an equal or lower cost, a smaller maximum observed arrival queue length, or both.

The motivation for these two features is to improve how SAGHP models represent the objectives of the decision to assign ground delay as it is made in practice. First, flight delay costs are assumed to be a non-linear function of the duration of ground, airborne, or cumulative delay, where cumulative delay refers to the sum of ground and airborne delay. Non-linear functions may more accurately represent costs such as flight diversions and missed connections than functions that are strictly linear. Previous models in the SAGHP literature that assign ground delays subject to a stochastic and dynamic forecasts of the airport arrival capacity do not consider non-linear cumulative or airborne delay costs.

Second, the ECM assumes that a possible objective of the design of a GDP in practice is to avoid lengthy airborne queues. SAGHP models that assign ground delays in order to minimize the total cost of delays to flights may permit arrival queues with lengths that would be unacceptable in practice. The ECM allows the traffic manager to specify a maximum limit, or capacity, of the airborne arrival queue permitted by the optimal solution.

After presenting the formulation of the model, the performance of the ECM is explored by two methods. One is an analysis of the sensitivity of the solution to the ECM to changes in the model's input parameters. The analysis identifies a relationship between these objectives and the times at which ground delays are assigned. As shown for a simple, hypothetical example, the total expected cost of the solution to the ECM is much less sensitive to the time at which ground delays are initially assigned subject to an uncertain arrival capacity forecast than to the time at which they are revised subject to perfect information. On the other hand, ability of the model to identify solutions with smaller arrival queues is much more sensitive of the time of the initial GDP than to subsequent revisions.

The second method is an arrival study case study analysis, which solves the SAGHP for a variety of arrival capacity forecasts. This analysis shows that not only can the ECM be applied to a wide variety of arrival capacity forecasts, but that it also results in solutions that are preferable to those of previous models in the literature. For each of the examples, the solutions to the ECM either results in an equal or lower cost, a smaller maximum observed arrival queue length, or both.

As an extension of the SAGHP, this dissertation also examines the design of a GDP in practice and the salience of various stakeholder groups. Various stakeholder groups are described and classified according to their power, legitimacy, and urgency in regards to the design of a GDP. Airlines (ASPs) and the FAA are identified as being the most salient stakeholders, which is in accordance with their roles in the GDP design process. However, passengers, are an important, but under-represented stakeholder in the design process. Although this result may not seem surprising (to passengers), previous SAGHP models have not considered how delays might affect passengers.

This dissertation presents a final analysis to demonstrate how passengers might be considered during the design of a GDP. A second MIP formulation is proposed assigns ground delays to minimize the total passenger delay cost. This formulation is applied to a simple example and the results show that the delay cost to passengers could be reduced by 35% if the initial assignment of ground delay were to explicitly consider passengers as opposed to a model that only considers delay costs to aircraft.

Throughout this dissertation, various problems are identified as questions for future research. These problems can be grouped into four general areas. First, research could explore the properties of the ECM as a SAGHP model. Further analyses could be conducted to explore the sensitivity of the solution to the simultaneous variation of multiple parameters. Such an analysis would provide insight not only for the design of GDPs in practice, but also into the time required by the ECM to find a solution. As highlighted by Chapters Four and Five, the ECM often, but not always, yields an integer solution when solved as an LP. Such solutions are preferable because the run time of the model is often significantly shorter if it can be solved as an LP. Therefore, additional research could seek to identify parameter values that are more or less likely to yield integer solutions or to identify an alternate formulation that has a shorter run time.

A second area of research could explore how the ECM might be used for the design of a Ground Delay Program in practice. In particular, the ECM requires a set of delay cost functions and an arrival capacity scenario. Further research could explore the development of these input parameters from practical data sources. A second area related to the application of the ECM would be to compare the solutions provide by the ECM for those of an actual GDP.

The third area for future research would be to explore the incorporation of additional considerations into the SAGHP. Other GHP models have looked to incorporate flight banks, the connections made by aircraft, and the capacity of en route sectors, etc. However, these models have generally considered deterministic or static arrival capacities. Thus, one area would be to explore these concepts under a stochastic and dynamic forecast and to determine their effects on the design of a GDP.

A fourth area of research is to further explore the role and relevance of various stakeholders. The analysis presented in Chapter Six identifies that certain groups have more salience than others. However, groups such as airlines contain many members, each of which may play a different role. Thus, a new stakeholder analysis could explore the salience of different airlines, passenger market segments, airports, communities, etc.

One area in which an understanding of the various stakeholders may be of special importance is the concept of equity as applied to a GDP. Initial research in this area has defined equity in regards to the equality of aircraft. However, the notion of equity does not necessarily imply equality among aircraft and may also need to consider additional stakeholder groups, as well. As an area for future work, identifying a more complete definition of equity is of particular interest to the author of this dissertation.

Appendices

Appendix 1: Acronyms and Glossary

AAM	Airline Assignment Model
ACID	Aircraft ID
ADC	Airborne Delay Cost
AOC	Airline Operations Center
ARTCC	Air Route Traffic Control Center (FAA facility)
ASP	Air Transportation Service Provider
ATC	Air Traffic Control
ATCSCC	Air Traffic Control System Command Center (FAA facility)
ATF	Aviation Trust Fund
ATFM	Air Traffic Flow Management
ATM	Air Traffic Management
CDC	Cumulative Delay Cost
CDM	Collaborative Decision Making
DP	Dynamic Programming
ECM	Extended Cost Model
ETMS	Enhanced Traffic Management System
FAA	Federal Aviation Administration
FCFS	First Come First Served (order)
FSFS	First Scheduled First Served (order)
FDCM	Flight-Defined Cost Model
GA	General Aviation
GDC	Ground Delay Cost

GDP	Ground Delay Program
GHG	Greenhouse Gas
GHP	Ground Hold Problem
IP	Integer Program
LP	Linear Program
M-DM	Mukherjee Dynamic Model
M-DMDD	Mukherjee Dynamic Model with Distributed Delays
MIP	Mixed-Integer Program
MIT	Miles-In-Trail
NAS	National Airspace System
OAG	Official Airline Guide
ORD	Chicago O'Hare International Airport
PAAR	Planned Arrival Acceptance Rate
PFC	Passenger Facility Charge
RBS	Ration-by-Schedule
RO-PDM	Richetta-Odoni Partially Dynamic Model
SAGHP	Single Airport Ground Hold Problem
SCS	Slot-Credit-Substitution
SFO	San Francisco International Airport
TM	Traffic Manager
TMU	Traffic Management Unit
TRACON	Terminal Radar Approach Control (FAA facility)
TSA	Transportation Security Administration

Cost under a profile: The cost of the solution to a model as evaluated by the set of common cost functions in Figure 3-08 and assuming the realization of a particular arrival capacity profile

Departure queue: At a given point in time, the number of aircraft being delayed by a GDP past their scheduled departure time

Irrelevant: A profile with a negligible likelihood

Marginal cost or marginal delay: The observed incremental change in cost or delay per unit of change in the independent variable with respect to the previous trial

Maximum observed arrival queue: The maximum number of flights that would simultaneously experience airborne delay under any profile

Maximum observed flight delay: The greatest amount of delay assigned to any one flight under any profile

Profile: A set of rates that represent a time-vary airport arrival capacity forecast

Realize: The occurrence of a particular profile

Scenario: A set of mutually exclusive and collectively exhaustive probabilistic arrival capacity profiles, each with an associated likelihood of occurring

Scenario Tree: The set of arrival capacity scenarios that may occur over time

Time of Perfect Information: The time at which the actual airport arrival capacity, or the profile that will be realized, is known with certainty

Total expected cost: The sum of the costs under each profile weighted by the profile likelihoods

Appendix 2: Model Formulations

Appendix 2.1: ECM and FDCM Notation

Notation

$b \in B$	Arrival capacity scenarios
$q \in Q$	Arrival capacity profiles
$t \in T$	Period of times
$f \in F$	Scheduled Flight Arrivals
p_q	The likelihood of capacity profile q
M_{qt}	The arrival capacity of the airport under profile q at time t
W_{qt}^{MAX}	The maximum allowable arrival queue under profile q at time t
b_s	The start time of scenario b
b_e	The end time of scenario b
α_i	The cost of the i^{th} period of airborne delay
γ_i	The cost of the i^{th} period of ground delay
χ_i	The cost of the i^{th} period of cumulative delay
α_i^f	The cost of the i^{th} period of airborne delay of flight f
γ_i^f	The cost of the i^{th} period of ground delay of flight f
χ_i^f	The cost of the i^{th} period of cumulative delay of flight f
STE_f	The scheduled time en route of flight f
$SAT_f \in T$	The scheduled arrival time of flight f
$XMT_f \in \{0,1\}$	1 for all exempt flights, 0 otherwise

Decision Variables

$d_{ftq} \in (0,1)$	an indicator if flight f has departed by time t in profile q
$\delta_{ftq} \in (0,1)$	an indicator if flight f has demanded arrival by time t in profile q
$\lambda_{ftq} \in (0,1)$	an indicator if flight f has landed by time t in profile q
Z_{fstq}	Helper variable to compute airborne delay costs of flight f in profile q

Appendix 2.2: ECM MIP Formulation

$$\textit{Minimize} \quad \sum_{Q}\left[p_q \times \sum_{F}\left(CDC_{fq} + ADC_{fq} + GDC_{fq}\right)\right] \tag{3.01}$$

$$CDC_{fq} = \sum_{t=SAT_f}^{T-1}\left(\left(1-\lambda_{ftq}\right)\times \chi_{t-SAT_f+1}\right) \qquad \forall\ f \in F;\ q \in Q \tag{3.02}$$

$$GDC_{fq} = \sum_{t=SDT_f}^{T+SDT_f-SAT_f-1}\left(\left(1-d_{ftq}\right)\times \gamma_{t-SDT_f+1}\right) \qquad \forall\ f \in F;\ q \in Q \tag{3.03}$$

$$ADC_{fq} = \sum_{t=SAT_f}^{T-1}\left(\left(1-\lambda_{ftq}\right)\times \alpha_{t-SAT_f+1}\right) - \sum_{s=0}^{T-SAT_f-1}\ \sum_{t=s}^{T-SAT_f-1} Z_{fstq} \qquad \forall\ f \in F;\ q \in Q \tag{3.04}$$

$$Z_{fstq} \le \left(\alpha_{t-s+1} - \alpha_{t-s}\right)\times\left(1-\lambda_{f\left(t+SAT_f\right)q}\right) \tag{3.06a}$$

$$Z_{fstq} \le \left(\alpha_{t-s+1} - \alpha_{t-s}\right)\times\left(1-d_{f\left(s+SDT_f\right)q}\right) \tag{3.06b}$$

$$\forall\ f \in F;\ s \in 0\ldots T-SAT_f-1;\ t \in s\ldots T-SAT_f-1;\ q \in Q$$

$$d_{ftq} \ge d_{f(t-1)q} \qquad \forall\ f \in F;\ t \in 1\ldots T;\ q \in Q \tag{3.07}$$

$$\lambda_{ftq} \ge \lambda_{f(t-1)q} \qquad \forall\ f \in F;\ t \in 1\ldots T;\ q \in Q \tag{3.08}$$

$$d_{f\left(t-STE_f\right)q} = \delta_{ftq} \qquad \forall\ f \in F;\ t \in STE_f\ldots T;\ q \in Q \tag{3.09}$$

$$\lambda_{ftq} \le \delta_{ftq} \qquad \forall\ f \in F;\ t \in \tau;\ q \in Q \tag{3.10}$$

$$\delta_{ftq} = 0 \qquad \forall\ f \in F;\ t \in 1\ldots SAT_f - 1;\ q \in Q \tag{3.11}$$

$$\lambda_{fTq} = 1 \qquad \forall\ f \in F;\ q \in Q \tag{3.12}$$

$$\sum_{F}\left(\delta_{ftq} - \lambda_{ftq}\right) \le W_{tq}^{MAX} \qquad \forall\ t \in \tau;\ q \in Q \tag{3.13}$$

$$\sum_{F}\left(\lambda_{ftq} - \lambda_{f(t-1)q}\right) \le M_{tq} \qquad \forall\ t \in \tau;\ q \in Q \tag{3.14}$$

$$d_{ftq_i} = \ldots = d_{ftq_j} \qquad \forall\ f \in F;\ b \in B;\ q_i,\ldots,q_j \in b;\ t \in \left\{b_s \ldots b_e\right\} \tag{3.15}$$

$$d_{f0q} = XMT_f \qquad \forall\ f \in F;\ q \in Q \tag{3.16}$$

Appendix 2.3: A Proof of Binary Decision Variables

Claim: If there exists an optimal solution to the ECM with binary values for d_{ftq}, then there exists an optimal solution with binary values for δ_{ftq} and λ_{ftq}.

Proof: This proof is presented in two parts, one for each of the variables.

Part I: By constraints (3.09), δ_{ftq} are fully defined by d_{ftq} and a constant, STE_f. As d_{ftq} are binary for any solution by definition, δ_{ftq} will also be binary.

Part II: For the second part of the proof, it is helpful to prove the two following lemmas:

Lemma #1: There will exist an optimal solution with an integer number of flight arrivals in each time period.

Proof: First, as the cost of flight delay is non-negative for each incremental period of delay, an optimal solution will utilize arrival capacity if there is a flight that demands arrival. Second, as the arrival demand (by Part I) and arrival capacities are integer (by definition), the cumulative number of flight arrivals will be integer.

1. For time periods in which the arrival capacity is sufficient to meet demand, all eligible flights will land and the corresponding values of λ_{ftq} will be 1 (and λ_{ftq} will be 0 for all other flights) (Lemma #1).

Lemma #2: Flight arrivals during each time period t and under each profile q occur in order of decreasing marginal flight delay cost.

Proof: If there are two flights that demand arrival but the capacity to accommodate only one arrival, then the flight that would incur the greater cost if delayed in the air for an (additional) period will be assigned to arrive during time period t. If the flight with the lower marginal delay cost arrives, instead, then it would increase the total delay cost and the solution would not be optimal.

2. For time periods in which the arrival capacity is insufficient to meet demand the optimal solution will allocate available arrival capacity to those flights with the greatest marginal delay cost (Lemma #2). Therefore, a solution can always be found in which the arrival variables are binary.

Corollary: If there exists an optimal solution to the ECM and if the incremental flight delay cost parameters adhere to the following relationship:

$$\alpha_i + \chi_n \le \alpha_j + \chi_j \qquad \forall \quad n;\ i < j \qquad \text{(AP1.01)}$$

then there exists an optimal solution in which flight arrivals are first-come-first serve (FCFS).

Proof: By the previous proof, there exists an optimal solution to the ECM for which an integer number of flights arrive during each time period (and under each profile) and the flights that arrive during each time period are those with the greatest marginal airborne delay cost. This solution will be consistent with a FCFS arrival process if those flights with the greatest marginal airborne delay cost are also those flights that entered the arrival queue earliest.

The marginal cost of flight delay to an aircraft is comprised of two cost components, airborne delay cost and cumulative delay cost. Considering only airborne delay cost, as the airborne delay cost function is both identical for all aircraft and non-decreasing in the duration of delay, those flights that have been in the arrival queue the longest are those that will incur the greater marginal *airborne* delay cost.

If cumulative delay costs are include, a complication might arise if one flight has a higher cumulative delay cost than another flight. If the additional cumulative delay is greater than the difference in airborne delay cost, then the flights might not arrive in FCFS order. By ensuring the relationship in AP1.01, costs due to additional cumulative delay will be less than the difference in airborne delay cost and the optimal solution will feature FCFS flight arrivals.

Appendix 2.3: The FDCM MIP Formulation

$$\textit{Minimize} \quad \sum_{Q}\left[p_q \times \sum_{F}\left(CDC_{fq} + ADC_{fq} + GDC_{fq}\right)\right] \tag{7.08}$$

$$CDC_{fq} = \sum_{t=SAT_f}^{T-1}\left(\left(1-\lambda_{ftq}\right) \times \chi^{f}_{t-SAT_f+1}\right) \qquad \forall\ f \in F;\ q \in Q \tag{7.09}$$

$$GDC_{fq} = \sum_{t=SDT_f}^{T+SDT_f-SAT_f-1}\left(\left(1-d_{ftq}\right) \times \gamma^{f}_{t-SDT_f+1}\right) \qquad \forall\ f \in F;\ q \in Q \tag{7.10}$$

$$ADC_{fq} = \sum_{t=SAT_f}^{T-1}\left(\left(1-\lambda_{ftq}\right) \times \alpha^{f}_{t-SAT_f+1}\right) - \sum_{s=0}^{T-SAT_f-1}\ \sum_{t=s}^{T-SAT_f-1} Z_{fstq} \qquad \forall\ f \in F;\ q \in Q \tag{7.11}$$

$$Z_{fstq} \le \left(\alpha^{f}_{t-s+1} - \alpha^{f}_{t-s}\right) \times \left(1-\lambda_{f(t+SAT_f)q}\right) \tag{7.12}$$

$$Z_{fstq} \le \left(\alpha^{f}_{t-s+1} - \alpha^{f}_{t-s}\right) \times \left(1-d_{f(s+SDT_f)q}\right) \tag{7.13}$$

$$\forall\ f \in F;\ s \in 0 \ldots T-SAT_f-1;\ t \in s \ldots T-SAT_f-1;\ q \in Q$$

$$d_{ftq} \ge d_{f(t-1)q} \qquad \forall\ f \in F;\ t \in 1 \ldots T;\ q \in Q \tag{7.14}$$

$$\lambda_{ftq} \ge \lambda_{f(t-1)q} \qquad \forall\ f \in F;\ t \in 1 \ldots T;\ q \in Q \tag{7.15}$$

$$d_{f(t-STE_f)q} = \delta_{ftq} \qquad \forall\ f \in F;\ t \in STE_f \ldots T;\ q \in Q \tag{7.16}$$

$$\lambda_{ftq} \le \delta_{ftq} \qquad \forall\ f \in F;\ t \in \tau;\ q \in Q \tag{7.17}$$

$$\delta_{ftq} = 0 \qquad \forall\ f \in F;\ t \in 1 \ldots SAT_f-1;\ q \in Q \tag{7.18}$$

$$\lambda_{fTq} = 1 \qquad \forall\ f \in F;\ q \in Q \tag{7.19}$$

$$\sum_{F}\left(\delta_{ftq} - \lambda_{ftq}\right) \le W^{MAX}_{tq} \qquad \forall\ t \in \tau;\ q \in Q \tag{7.20}$$

$$\sum_{F}\left(\lambda_{ftq} - \lambda_{f(t-1)q}\right) \le M_{tq} \qquad \forall\ t \in \tau;\ q \in Q \tag{7.21}$$

$$d_{ftq_i} = \ldots = d_{ftq_j} \qquad \forall\ f \in F;\ b \in B;\ q_i,\ldots,q_j \in b;\ t \in \{b_s \ldots b_e\} \tag{7.22}$$

$$\lambda_{ftq_i} = \ldots = \lambda_{ftq_j} \qquad \forall\ f \in F;\ b \in B;\ q_i,\ldots,q_j \in b;\ t \in \{b_s \ldots b_e\} \tag{7.23}$$

Appendix 2.4: The AAM LP Formulation and Notation

Notation

F_{ij} Equals 1 if flight i may be assigned to slot j and 0 otherwise

C_{ij} The cost of assigning flight i to slot j

Decision Variables

$X_{ij} \in \{0,1\}$ an indicator of whether flight i is assigned to slot j

Objective Function and Constraints

Minimize:
$$\sum_{I,J} C_{ij} \times X_{ij} \tag{7.03}$$

Subject To:
$$\sum_{J} X_{ij} = 1 \qquad \forall \; i \in I \tag{7.04}$$

$$\sum_{I} X_{ij} = 1 \qquad \forall \; j \in J \tag{7.05}$$

$$X_{ij} \leq F_{ij} \qquad \forall \; i \in I; \; j \in J \tag{7.06}$$

Appendix 2.5: Aircraft Seats as Defined by Aircraft Type

Aircraft Type	Count of Type	Seats
A318	1	107
A319	13	124
A320	16	150
A332	1	253
A333	1	335
A343	0	295
B721	0	149
B732	1	100
B733	24	130
B734	0	168
B735	10	132
B737	0	149
B738	1	189
B73C	0	189
B73F	1	189
B73J	0	189
B73S	0	189
B741	1	452
B742	1	452
B744	3	524
B747	0	467
B752	23	200
B757	0	243
B762	1	224
B763	6	269
B767	0	304
B772	8	305
BA46	9	50
CARJ	2	50
CRJ	1	50
CRJ1	15	50
CRJ2	44	50
CRJ7	21	70
DC85	0	200
DC9	0	150
DC93	2	115
DC94	1	125
E135	12	37
E145	55	50
E170	10	78
E45X	1	50
MD10	0	150
MD11	0	323
MD80	5	155
MD82	38	152
MD83	19	155
MD90	1	153
S319	0	50

"Aircraft Type" is taken from ETMS

"Count of Type" refers to the number of aircraft of each type that are included in the set of 348 flights used for examples in this thesis.

"Seats" refers to the total number of passenger seats in all classes in a common configuration of the aircraft.
Source: en.wikipedia.org, February 17, 2008
For cases in which there are multiple configurations, an average is used.

Appendix 3: Detailed Results for the Sensitivity Analysis

Figure AP3-01: A summary of the input parameters used for the experiments in the sensitivity analysis

					Decision Time		Incremental Delay Cost			
Exp.	M_{1t}	M_{2t}	p_1	p_2	Initial	Revision	Airborne	Ground	Cumulative	W_{qt}^{MAX}
#1	25	***Varies***	**0.10**	**0.90**	1300 Z	1500 Z	α_t	γ_t	χ_t	30
#2	25	0	***Varies***	***Varies***	1300 Z	1500 Z	α_t	γ_t	χ_t	30
#3	25	0	**0.95**	**0.05**	1300 Z	1500 Z	α_t	γ_t	χ_t	***Varies***
#4	25	0	0.90	0.10	***Varies***	1500 Z	α_t	γ_t	χ_t	**348**
#5	25	0	0.90	0.10	***Varies***	1500 Z	α_t	γ_t	χ_t	***Varies***
#6	25	0	0.90	0.10	1300 Z	***Varies***	α_t	γ_t	χ_t	30
#7	25	0	**0.75**	**0.25**	1300 Z	1500 Z	$\boldsymbol{\alpha_t'}$	γ_t	χ_t	**348**
#8	25	0	**0.75**	**0.25**	1300 Z	1500 Z	$\boldsymbol{\alpha_t''}$	γ_t	χ_t	**348**
#9	25	0	**0.75**	**0.25**	1300 Z	1500 Z	$\boldsymbol{\alpha_t'''}$	γ_t	χ_t	**348**

The following pages contain the detailed experimental results.

Units for Figures

Max. Queue:	aircraft	Run Time:	seconds
Exp. Cost, Cost:	units of cost	p_2:	likelihood of Profile #2
Air, Ground Delay:	aircraft-periods	∂:	units of cost / change in p_2

* Note, due to size, detailed results for experiment #5 are not reported in this appendix

Figure AP3-02: Results from experiment #1

Experiment #1: Arrival Capacity Under Profile #2

M_{2t}	Exp. Cost	∂	Profile #1 (0.10) Cost	Air Delay	Ground Delay	Max. Queue	Profile #2 (0.90) Cost	Air Delay	Ground Delay	Max. Queue	Run Time
0	1,157.2		349	0	349	0	1,247	55	995	16	8.71
5	697.7	91.9	245	0	245	0	748	8	732	5	9.07
10	412.0	57.1	97	0	97	0	447	0	447	0	8.35
15	165.3	49.3	24	0	24	0	181	0	181	0	6.63
17	95.8	34.8	13	0	13	0	105	0	105	0	5.34
18	65.1	30.7	12	0	12	0	71	0	71	0	5.25
20	29.8	17.7	10	0	10	0	32	0	32	0	4.67
22	20.8	4.5	10	0	10	0	22	0	22	0	4.34
24	12.7	4.1	10	0	10	0	13	0	13	0	4.32
25	10.0	2.7	10	0	10	0	10	0	10	0	4.22
26	9.1	0.9	10	0	10	0	9	0	9	0	4.17
28	7.3	0.9	10	0	10	0	7	0	7	0	4.21
29	6.4	0.9	10	0	10	0	6	0	6	0	4.18
30	6.4	0.0	10	0	10	0	6	0	6	0	4.19

Figure AP3-03: Results from experiment #2

Experiment #2: Profile Likelihood

p_2	Exp. Cost	∂^2	Profile #1 Cost	Air Delay	Ground Delay	Max. Queue	Profile #2 Cost	Air Delay	Ground Delay	Max. Queue	Run Time
0.00	160.0		160	0	160	0	1,581	151	899	30	9.26
0.05	229.8	-62	160	0	160	0	1,556	134	916	30	6.92
0.10	296.5	-99	173	0	173	0	1,408	116	934	30	7.15
0.15	358.3	-64	173	0	173	0	1,408	116	934	30	7.35
0.20	416.8	-42	189	0	189	0	1,328	100	950	30	7.40
0.25	473.3	-16	191	0	191	0	1,320	98	952	30	7.39
0.30	528.9	0	195	0	195	0	1,308	94	956	30	7.70
0.35	584.6	0	195	0	195	0	1,308	94	956	30	7.72
0.40	640.2	0	195	0	195	0	1,308	94	956	30	7.65
0.45	695.9	0	195	0	195	0	1,308	94	956	30	7.90
0.50	751.5	-14	195	0	195	0	1,308	94	956	30	8.12
0.55	806.5	-52	213	0	213	0	1,292	86	964	28	8.15
0.60	858.8	-24	233	0	233	0	1,276	78	972	26	7.98
0.65	910.0	0	245	0	245	0	1,268	74	976	25	7.97
0.70	961.1	-15	245	0	245	0	1,268	74	976	25	8.05
0.75	1,011.5	-2	257	0	257	0	1,263	71	979	24	8.20
0.80	1,061.8	-4	257	0	257	0	1,263	71	979	24	8.30
0.85	1,111.9	-96	301	0	301	0	1,255	63	987	20	8.09
0.90	1,157.2	-57	349	0	349	0	1,247	55	995	16	8.70
0.95	1,199.7	-42	433	0	433	0	1,240	48	1,002	16	9.80
1.00	1,240.0		433	0	433	0	1,240	48	1,002	16	10.18

Figure AP3-04: Results from experiment #3

Experiment #3: Capacity of the Arrival Queue

		Profile #1 (0.95)				Profile #2 (0.05)				
W^{MAX}	Exp. Cost	Cost	Air Delay	Ground Delay	Max. Queue	Cost	Air Delay	Ground Delay	Max. Queue	Run Time
60	106.2	10	0	10	0	1,934	246	804	56	5.71
56	106.2	10	0	10	0	1,934	246	804	56	5.76
52	109.6	15	0	15	0	1,907	239	811	52	5.64
48	115.4	23	0	23	0	1,871	227	823	48	5.72
44	127.1	38	0	38	0	1,820	210	840	44	6.40
40	150.9	66	0	66	0	1,764	186	864	40	6.96
38	164.8	82	0	82	0	1,738	176	874	38	7.17
36	179.5	100	0	100	0	1,690	164	886	36	6.95
34	196.2	120	0	120	0	1,644	154	896	34	6.79
32	212.9	140	0	140	0	1,598	144	906	32	6.94
30	229.8	160	0	160	0	1,556	134	916	30	6.87
28	247.6	181	0	181	0	1,513	121	929	28	7.05
26	267.8	205	0	205	0	1,460	106	944	26	7.18
24	292.1	233	0	233	0	1,415	95	955	24	7.29
22	316.6	261	0	261	0	1,373	85	965	22	7.08
20	341.1	289	0	289	0	1,331	75	975	20	7.03
18	367.5	319	0	319	0	1,289	65	985	18	7.22
16	393.9	349	0	349	0	1,247	55	995	16	7.41
15	Infeasible									

Figure AP3-05: Results from experiment #4

Experiment #4: Time of the Initial Decision to Implement a Program

			Profile #1 (0.90)				Profile #2 (0.10)				
Time	Exp. Cost	Cum. Δ%	Cost	Air Delay	Ground Delay	Max. Queue	Cost	Air Delay	Ground Delay	Max. Queue	Run Time
1300 Z	200.4		25	0	25	0	1,779	225	825	54	5.72
1315 Z	200.6	0.1%	24	0	24	0	1,790	226	824	54	5.35
1330 Z	200.8	0.2%	23	0	23	0	1,801	227	823	54	5.02
1345 Z	201.6	0.6%	19	0	19	0	1,845	231	819	54	4.71
1400 Z	201.8	0.7%	18	0	18	0	1,856	232	818	54	4.52
1415 Z	202.4	1.0%	10	0	10	0	1,934	246	804	56	4.08
1430 Z	203.4	1.5%	10	0	10	0	1,944	248	802	58	3.88
1445 Z	203.9	1.7%	10	0	10	0	1,949	249	801	59	3.65
1500 Z	205.1	2.3%	10	0	10	0	1,961	253	797	60	3.57

Figure AP3-06: Results from experiment #6

Experiment #6: Time of the Revision

T_{PI}	Exp. Cost	Δ%	Profile #1 (0.90) Cost	Air Delay	Ground Delay	Max. Queue	Profile #2 (0.10) Cost	Air Delay	Ground Delay	Max. Queue	Run Time
1300 Z	133.0		10	0	10	0	1,240	48	1,002	16	5.13
1315 Z	138.8	4%	15	0	15	0	1,253	53	997	18	5.12
1330 Z	152.6	10%	19	0	19	0	1,355	89	961	29	20.16
1345 Z	160.0	5%	23	0	23	0	1,393	87	963	30	5.40
1400 Z	172.6	8%	38	0	38	0	1,384	100	950	30	5.55
1415 Z	203.0	18%	71	0	71	0	1,391	99	951	30	5.85
1430 Z	243.8	20%	116	0	116	0	1,394	102	948	30	6.36
1445 Z	274.1	12%	149	0	149	0	1,400	108	942	30	6.65
1500 Z	296.5	8%	173	0	173	0	1,408	116	934	30	7.08
1515 Z	409.2	38%	299	0	299	0	1,401	109	941	30	7.89
1530 Z	565.4	38%	473	0	473	0	1,397	105	945	30	9.51
1545 Z	660.2	17%	577	0	577	0	1,409	117	933	30	10.43
1600 Z	664.1	1%	577	0	577	0	1,448	146	904	30	12.42

Figure AP3-07: Results from experiment #7

Experiment #7: Maximum Incremental Airborne Delay Cost

A	Exp. Cost	Profile #1 (0.75) Cost	Air Delay	Ground Delay	Max. Queue	Profile #2 (0.25) Cost	Air Delay	Ground Delay	Max. Queue	Run Time
1.00	331.5	12	0	12	0	1,290	240	810	56	4.03
2.00	368.5	17	0	17	0	1,423	230	820	55	4.37
3.00	401.5	60	0	60	0	1,426	178	872	46	4.44
4.00	411.5	65	0	65	0	1,451	172	878	45	5.17
5.00	421.0	71	0	71	0	1,471	165	885	44	5.22
6.00	422.8	71	0	71	0	1,478	165	885	44	6.04
7.00	424.5	71	0	71	0	1,485	165	885	44	6.20
8.00	426.3	71	0	71	0	1,492	165	885	44	6.25
9.00	428.0	71	0	71	0	1,499	165	885	44	6.49
10.00	429.8	71	0	71	0	1,506	165	885	44	6.12
15.00	438.5	71	0	71	0	1,541	165	885	44	6.65
25.00	456.0	71	0	71	0	1,611	165	885	44	7.22
35.00	473.5	71	0	71	0	1,681	165	885	44	7.12
45.00	491.0	71	0	71	0	1,751	165	885	44	6.95
55.00	508.5	71	0	71	0	1,821	165	885	44	7.10

Figure AP3-08: Results from experiment #8

Experiment #8: Incremental Airborne Delay Cost (Multiplicative)

		Profile #1 (0.75)				Profile #2 (0.25)				
B	Exp. Cost	Cost	Air Delay	Ground Delay	Max. Queue	Cost	Air Delay	Ground Delay	Max. Queue	Run Time
0.00	270.0	10	10	0	6	1,050	1,001	49	85	3.38
0.25	325.3	10	0	10	0	1,271	246	804	56	5.46
0.50	368.3	44	0	44	0	1,341	200	850	50	5.93
0.75	401.3	56	0	56	0	1,437	182	868	47	6.15
1.00	431.5	71	0	71	0	1,513	165	885	44	6.94
1.25	459.1	85	0	85	0	1,581	151	899	42	7.57
1.50	485.6	94	0	94	0	1,661	145	905	41	82.56
1.75	507.6	139	0	139	0	1,614	120	930	36	8.19
2.00	526.8	149	0	149	0	1,660	115	935	35	8.10
2.50	563.0	159	0	159	0	1,775	110	940	34	8.33
3.00	599.3	159	0	159	0	1,920	110	940	34	8.27
3.50	634.0	213	0	213	0	1,897	86	964	28	8.23
4.00	663.3	233	0	233	0	1,954	78	972	26	9.20
4.50	691.5	233	0	233	0	2,067	78	972	26	8.78
5.00	718.8	245	0	245	0	2,140	74	976	25	8.25
6.00	773.3	245	0	245	0	2,358	74	976	25	9.83
7.00	827.8	245	0	245	0	2,576	74	976	25	9.21
8.00	881.3	257	0	257	0	2,754	71	979	24	9.11
9.00	934.5	257	0	257	0	2,967	71	979	24	8.43
10.00	987.8	257	0	257	0	3,180	71	979	24	9.11

Figure AP4-09: Results from experiment #9

Experiment #9: Incremental Airborne Delay Cost (Additive)

		Profile #1 (0.75)				Profile #2 (0.25)				
C	Exp. Cost	Cost	Air Delay	Ground Delay	Max. Queue	Cost	Air Delay	Ground Delay	Max. Queue	Run Time
0.00	431.5	71	0	71	0	1,513	165	885	44	6.96
0.50	451.4	85	0	85	0	1,551	151	899	42	7.11
1.00	470.3	85	0	85	0	1,626	151	899	42	7.23
1.50	489.1	94	0	94	0	1,675	145	905	41	89.61
2.00	507.3	103	0	103	0	1,720	140	910	40	7.55
3.00	536.8	149	0	149	0	1,700	115	935	35	7.30
4.00	564.3	159	0	159	0	1,780	110	940	34	7.72
5.00	590.3	213	0	213	0	1,722	86	964	28	7.05
6.00	610.8	233	0	233	0	1,744	78	972	26	7.10
8.00	648.8	245	0	245	0	1,860	74	976	25	7.69
10.00	685.8	245	0	245	0	2,008	74	976	25	7.83
13.00	739.3	257	0	257	0	2,186	71	979	24	9.53
16.00	791.5	301	0	301	0	2,263	63	987	20	9.81
19.00	834.8	349	0	349	0	2,292	55	995	16	8.84
22.00	876.0	349	0	349	0	2,457	55	995	16	8.57
25.00	917.3	349	0	349	0	2,622	55	995	16	9.17

Appendix 4: Arrival Capacity Profiles and Scenario Trees

Figure AP4-01: Scenario for Case #1

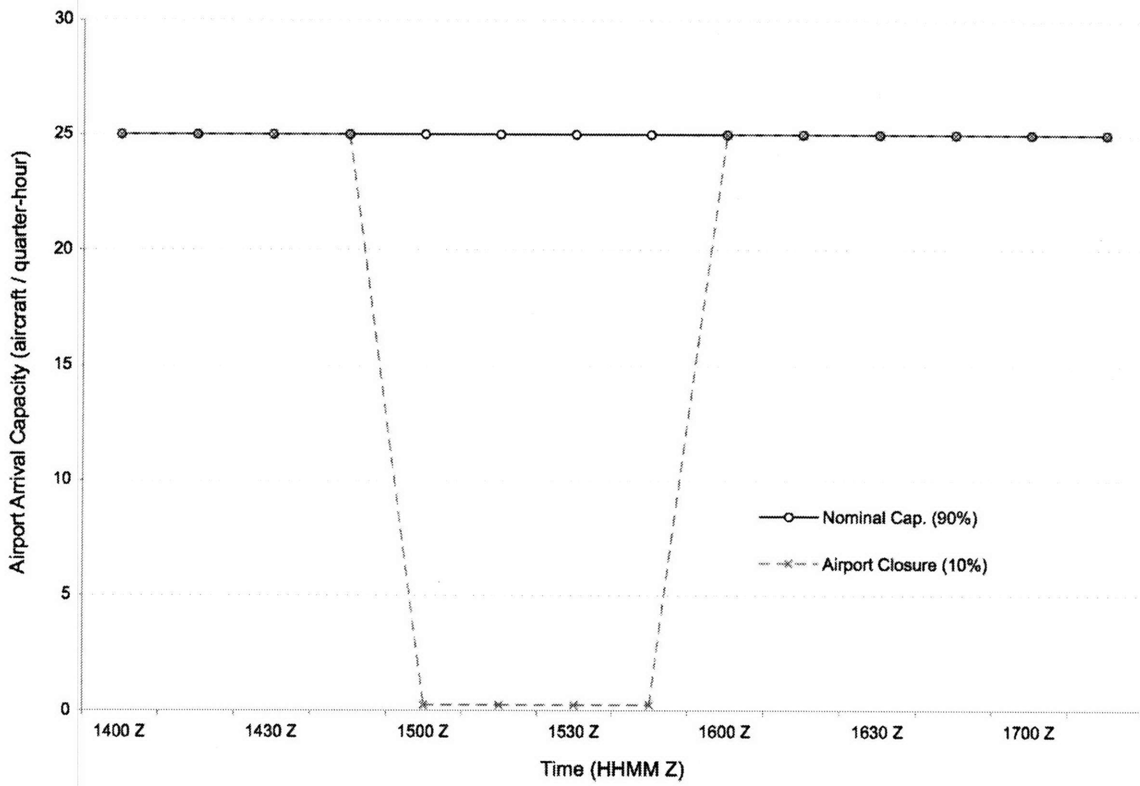

Figure AP4-02: Case #1 Scenario Tree

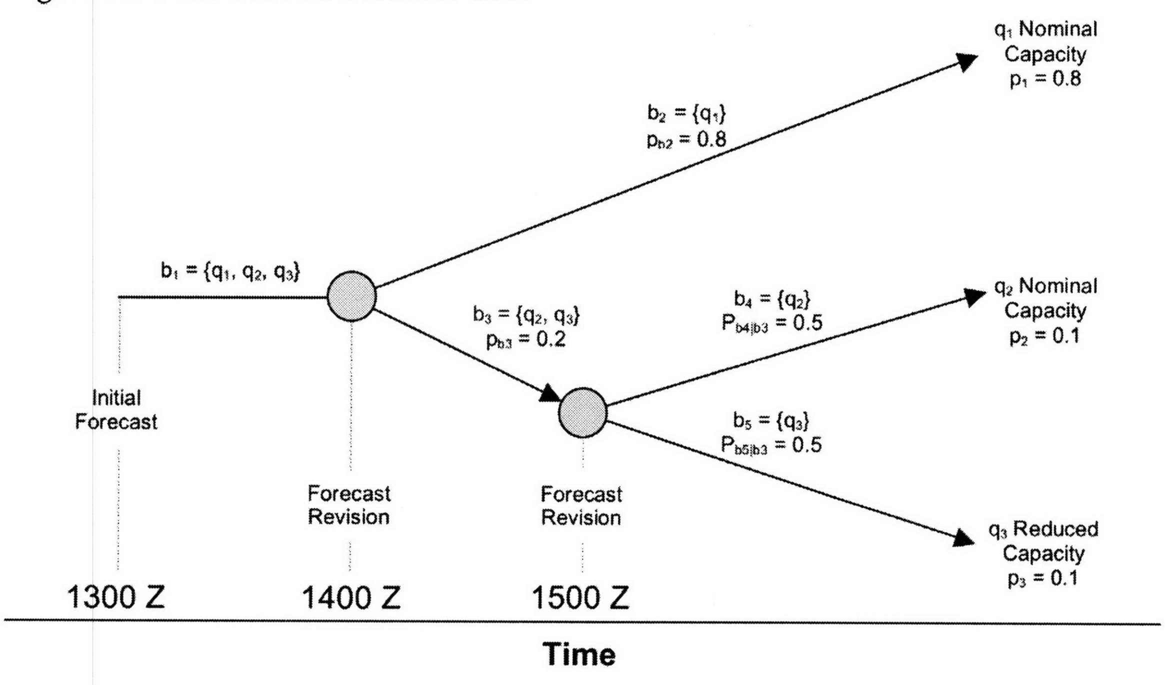

Figure AP4-03: Case #2 Arrival Capacity Profiles

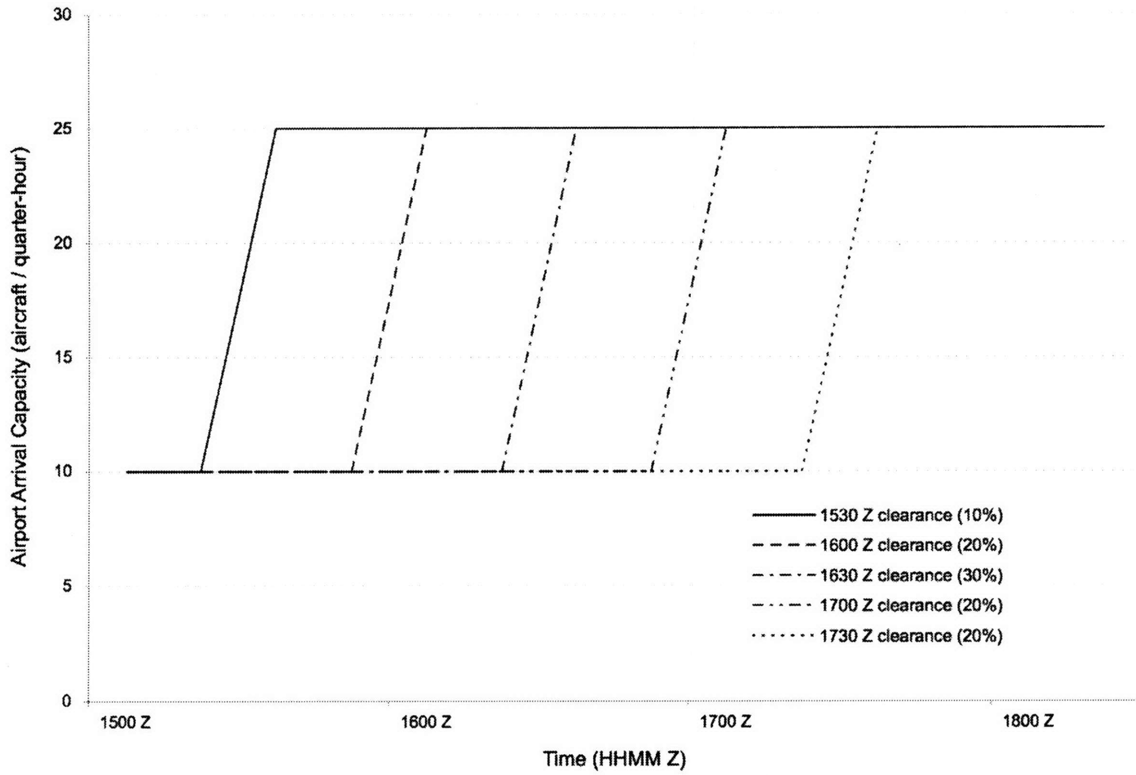

Figure AP4-04: Case #2 Arrival Capacity Scenario Tree

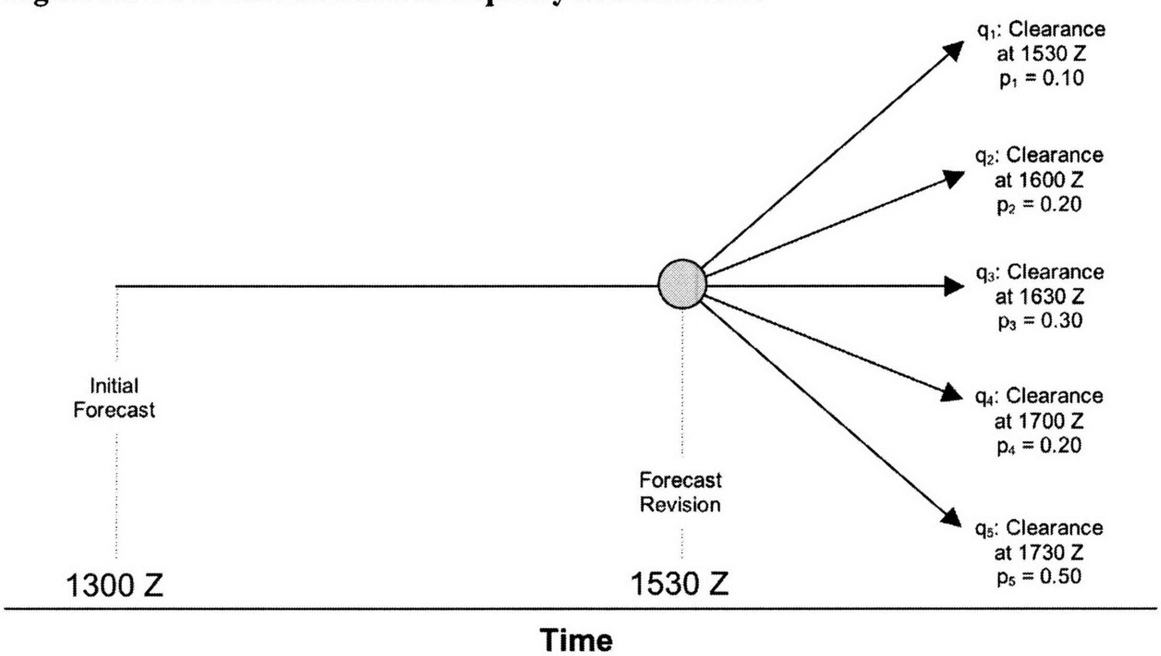

Figure AP4-05: Case #3 Scenario Tree

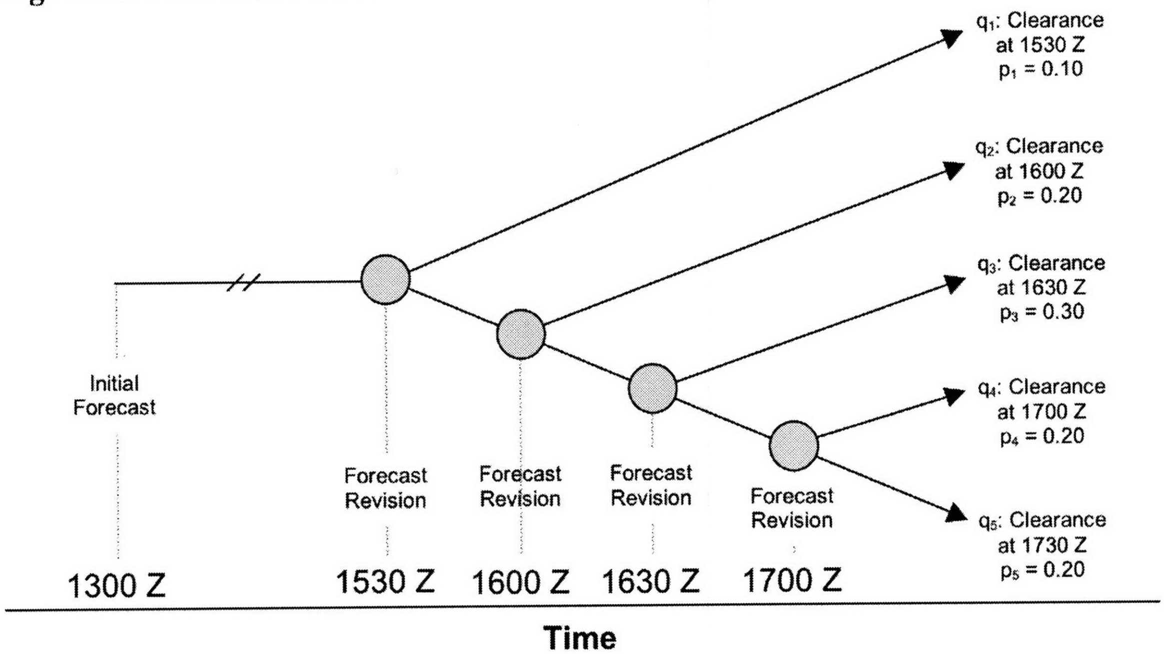

Figure AP4-06: Case #4 Scenario Tree

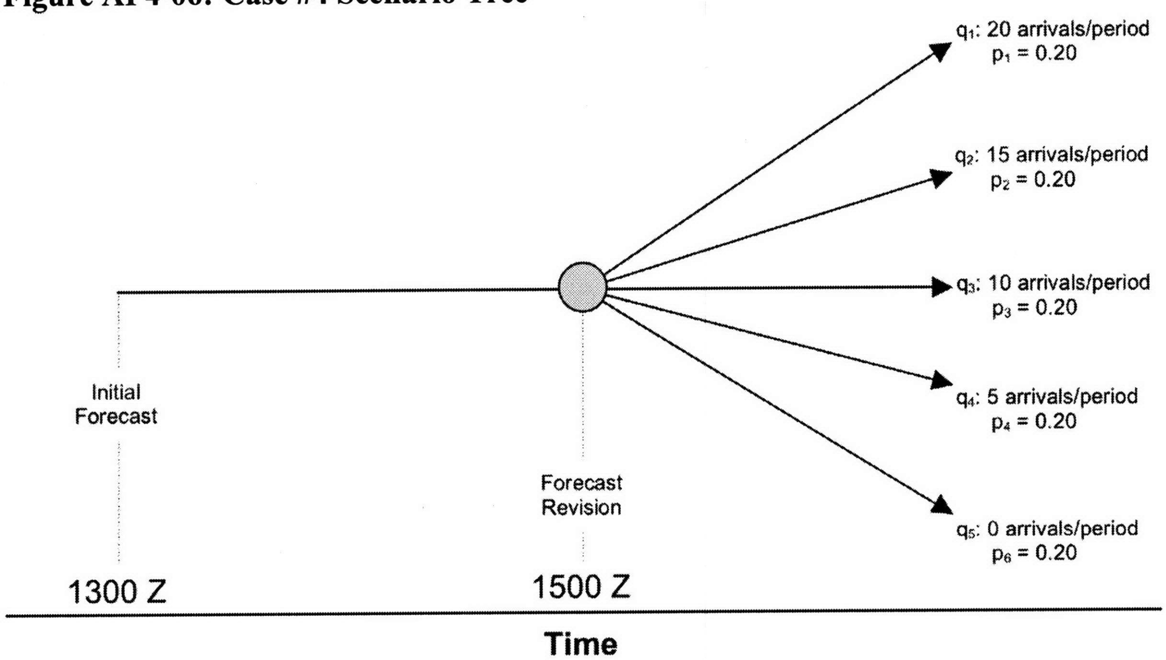

Figure AP4-07: Scenario for Cases #5/6

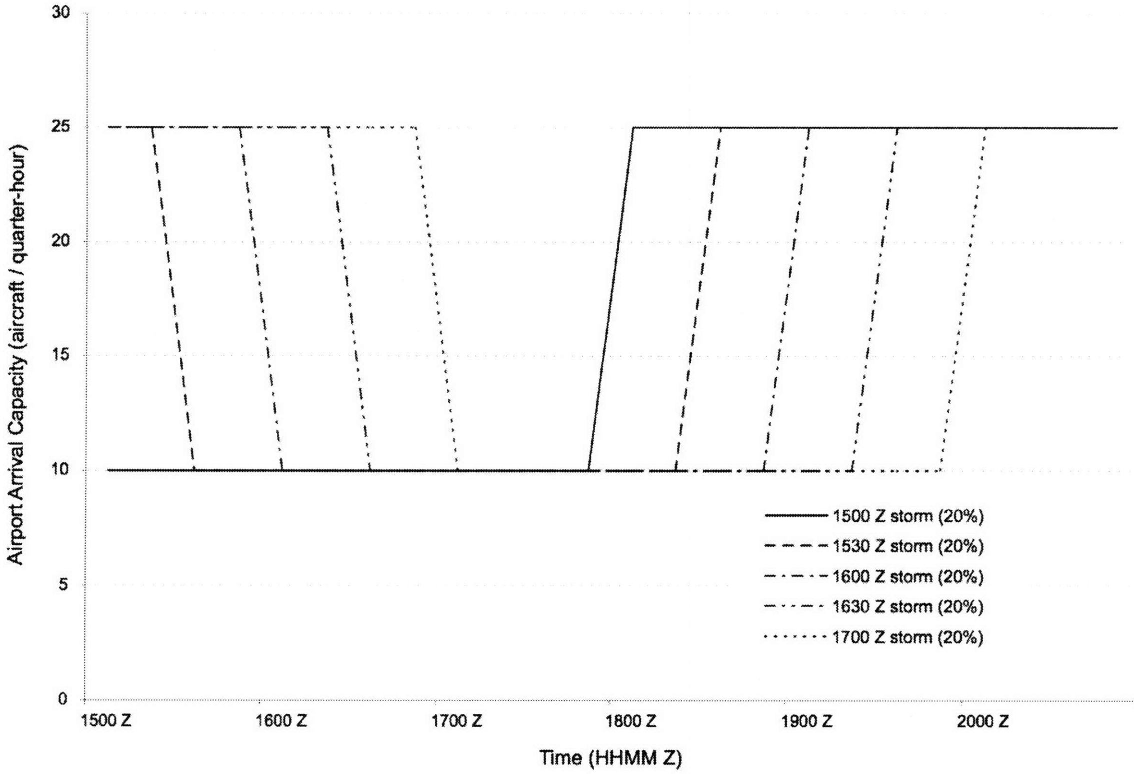

Figure AP4-08: Case #5 Scenario Tree

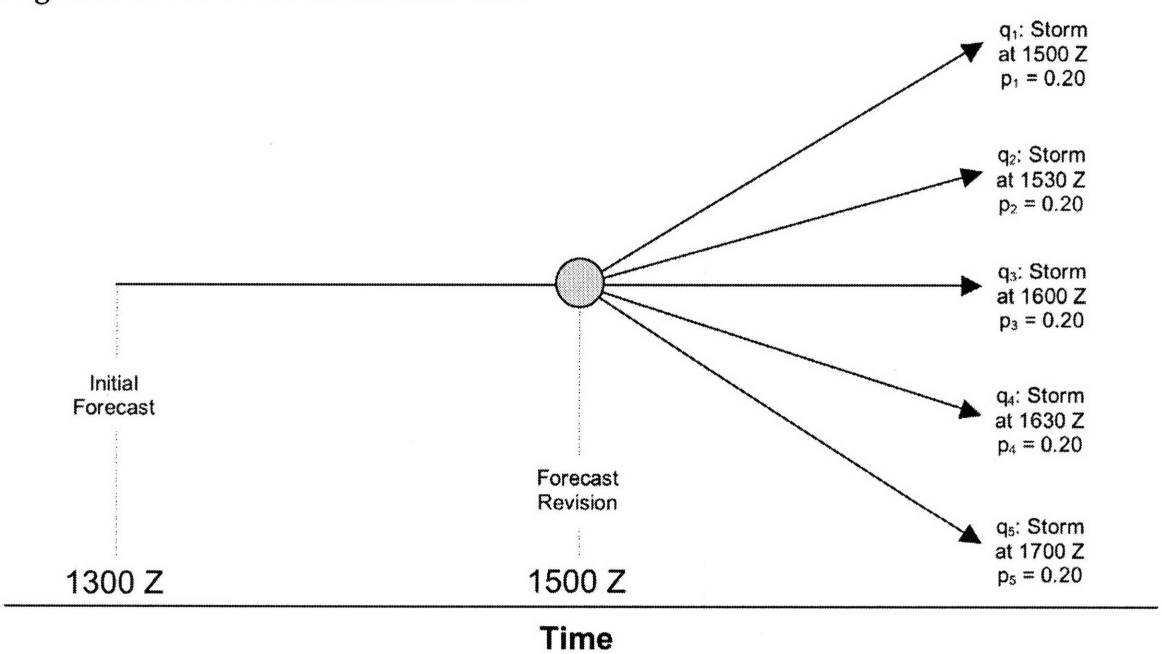

Figure AP4-09: Case #6 Scenario Tree

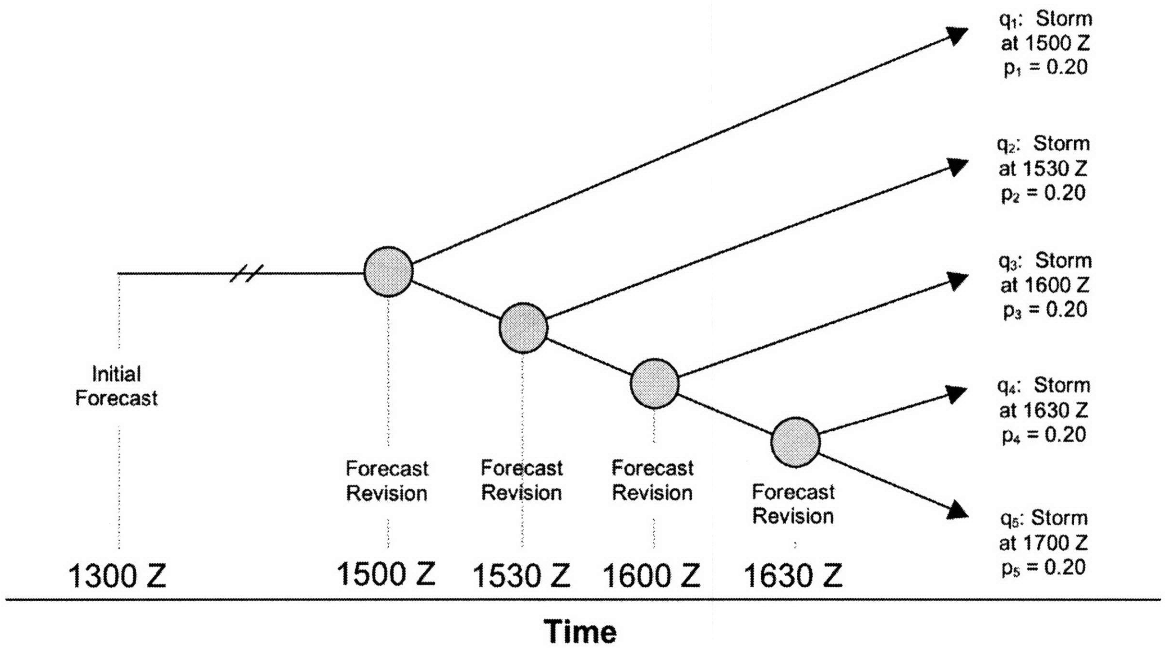

Figure AP4-10: Scenario for Cases #7/8

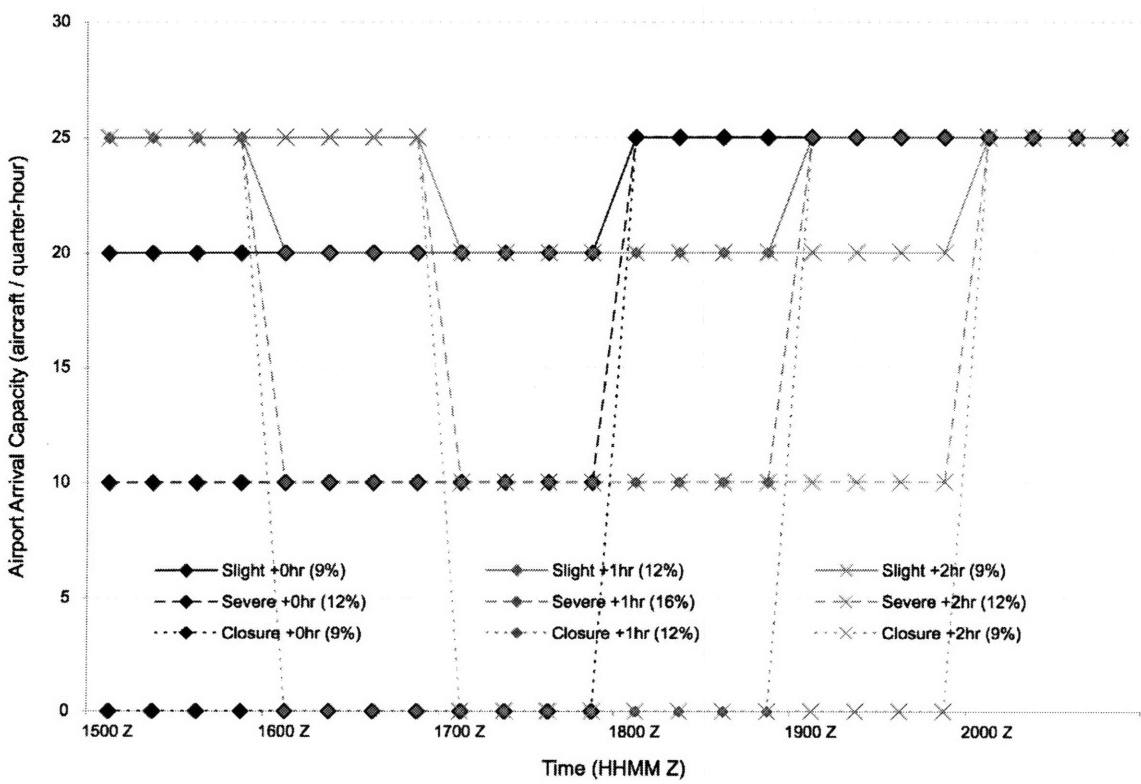

Figure AP4-11: Scenario Tree for Case #7

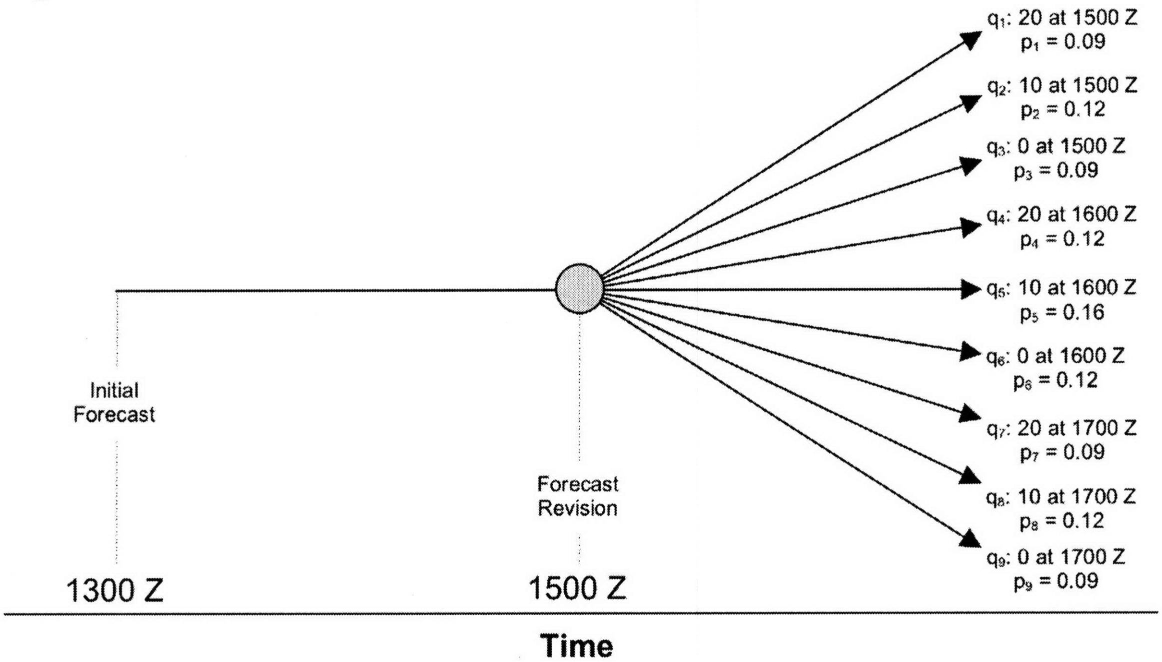

Figure AP4-12: Scenario Tree Case #8

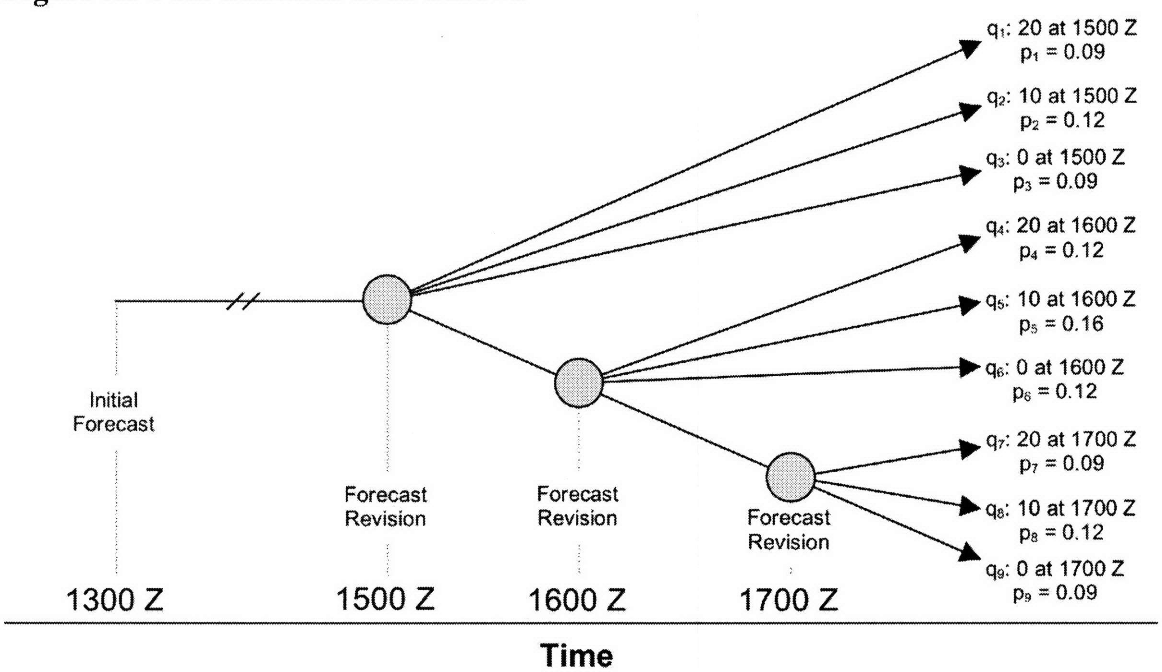

Figure AP4-13: Scenario for Cases #9/10

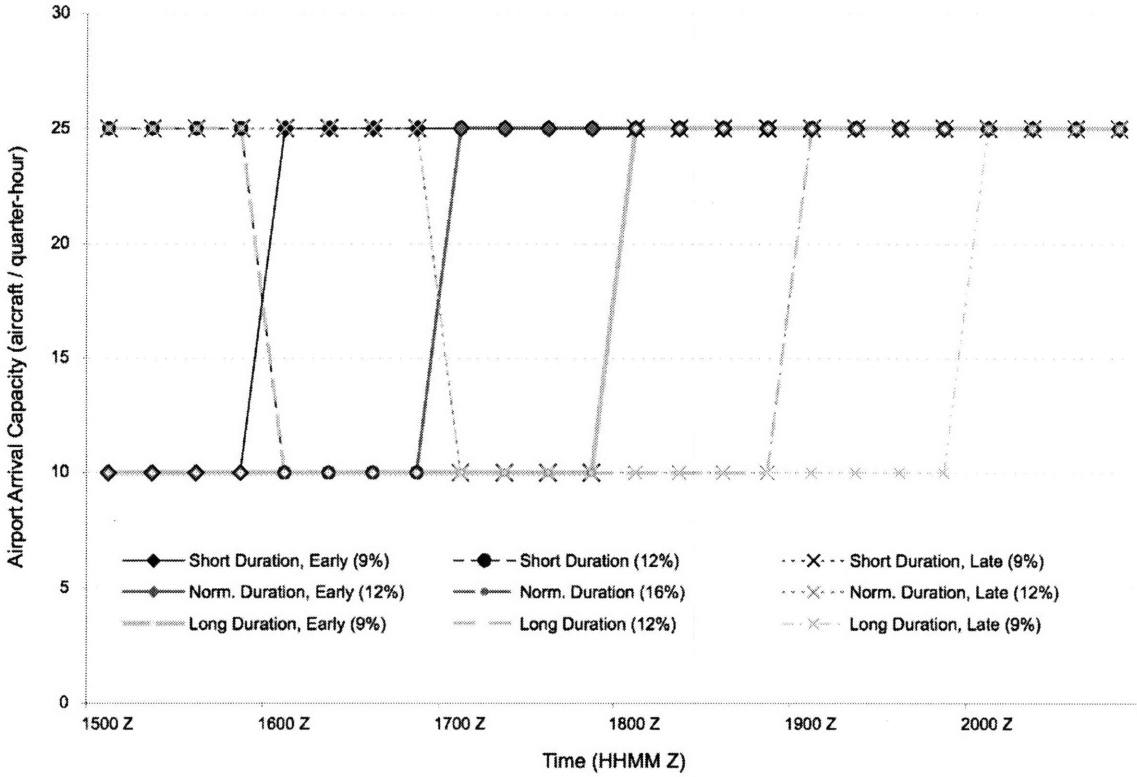

Figure AP4-14: Scenario Tree for Case #9

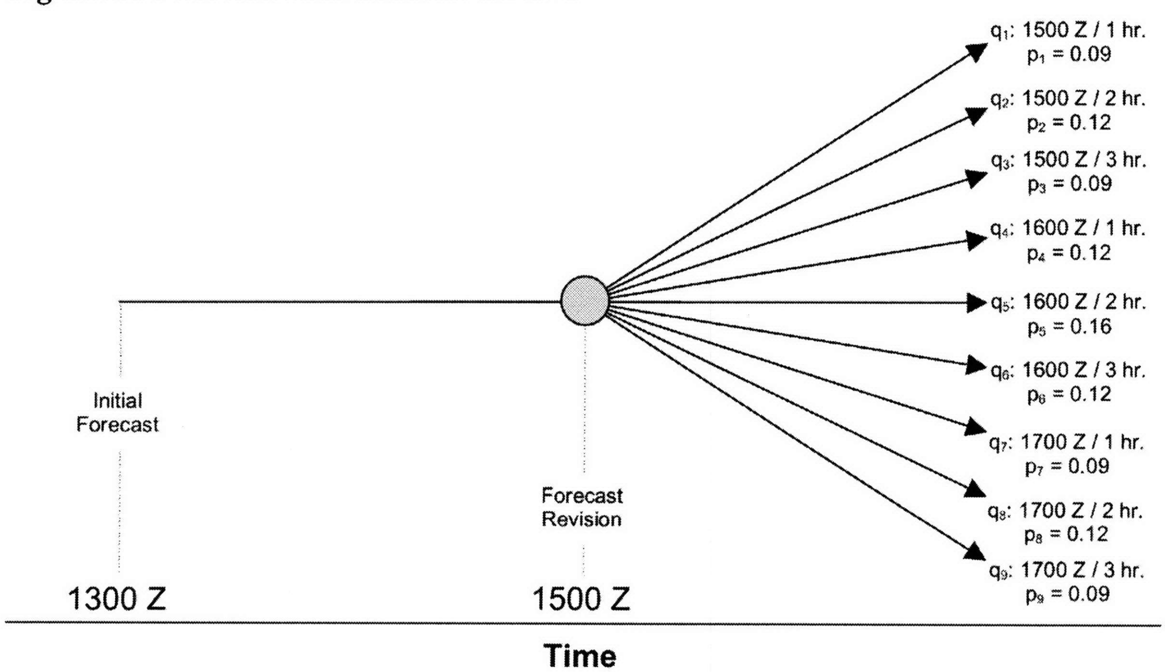

Figure AP4-15: Scenario Tree for Case #10

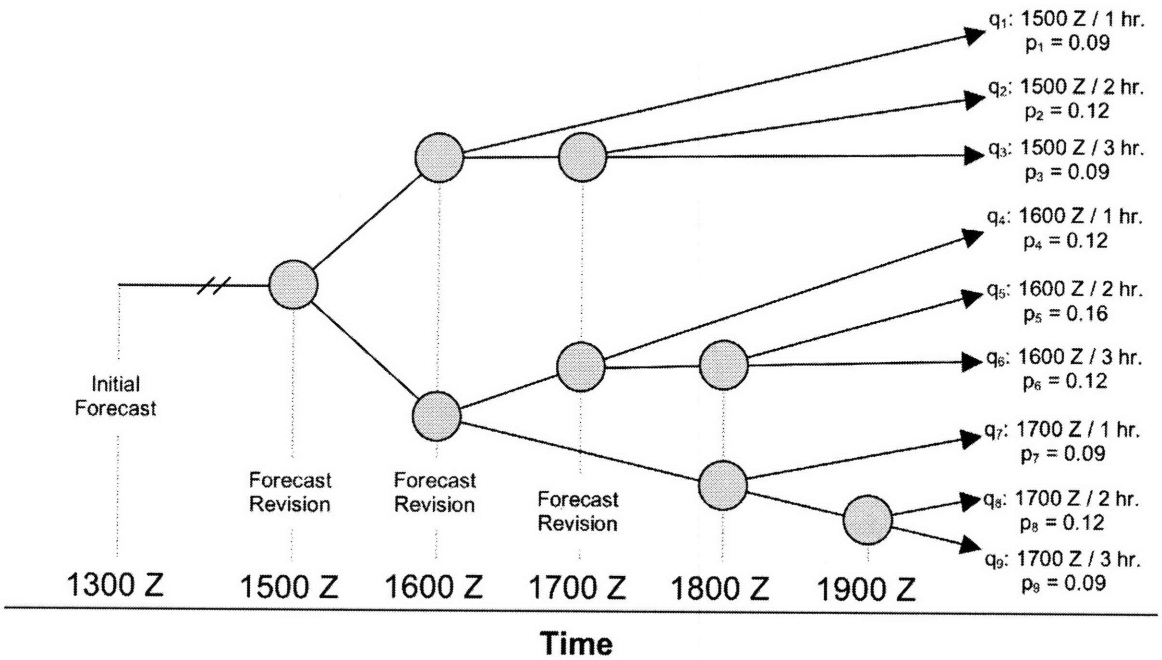

Appendix 5: Stakeholder Analysis Supplemental Figures

Figure AP5-01: Stakeholder relationships, Air Transportation Service Providers

Relationship	**Air Transportation Service Provider provides (to)**	**Air Transportation Service Provider receives (from)**
Airports	Rent and usage fees Infrastructure (and financing for construction) [1] Access to passengers	Essential air transport facility services/infrastructure Access to passengers
Business	Direct revenue Indirect revenue via: Transportation services Promise of future service Access to passengers	Miscellaneous business services and products Fares and service fees
FAA	Information/Feedback Info (e.g. flight plans, schedules, status) Information for use in Air Traffic Control (e.g. ride reports) Finance (via Aviation Trust Fund)	Industry regulation, certification, and oversight Essential flight services: Air Traffic Control Communication Navigation Airspace management
Federal Government	Taxes (via Treasury) Federal lobby (via lobbying organizations)	Threat of economic regulation Security (e.g. TSA) General services
Local/Regional/State Government		Tax and other incentives Employment base General services
Passengers	Transportation services Promise of future service Amenities and information Contractual obligation	Fares and service fees Information/Feedback (direct/indirect)
Society	Wages Environmental pollution (noise, particulates, GHG)	Labor

Figure AP5-02: Stakeholder relationships, Airports

Relationship	Airport(s) provides (to)	Airport(s) receives (from)
Air Transportation Service Providers (ASPs)	Essential air transport facility services/infrastructure Access to passengers	Rent and usage fees Infrastructure (and financing for construction) [1] Access to passengers
Business	Revenue Access to airlines Access to passengers Access to air transportation	Miscellaneous business services and products Revenue (rent, other)
FAA		Industry regulation, certification, and oversight Navigational services Grants in Aid (for safety, capacity, security) [2]
Federal Government	Taxes (via Treasury) Critical infrastructure (i.e. national security)	Threat of regulation and oversight (e.g. TSA) General services
Local/Regional/State Government	Revenue [3] (taxes, other)	Appointments (mandate) Local infrastructure Financing for construction
Passengers	Essential facility services Access to air transportation Amenities, concessions, and parking	Revenue (PFC, taxes, fees)
Society	Wages Environmental pollution (noise, particulates, GHG) Information (official reports)	Labor

Figure AP5-03: Stakeholder relationships, Business

Relationship	Business provides (to)	Business receives (from)
Air Transportation Service Providers (ASPs)	Miscellaneous business services and products Fares and service fees	Direct revenue Indirect revenue via: Transportation services Promise of future service Access to passengers
Airports	Miscellaneous business services and products Revenue (rent, other)	Revenue Access to airlines Access to passengers Access to air transportation
FAA	Miscellaneous business services and products	Revenue
Federal Government	Goods and services Taxes (via Treasury) Federal lobby (via lobbying organizations)	Revenue General services Threat of regulation and oversight
Local/Regional/State Government	Taxes Goods and services	General services Revenue Public (local) infrastructure
Passengers	Wages Goods and services	Labor Revenue
Society	Wages Goods and services	Labor Revenue

Figure AP5-04: Stakeholder relationships, FAA

Relationship	**FAA provides (to)**	**FAA receives (from)**
Air Transportation Service Providers (ASPs)	Industry regulation, certification, and oversight Essential flight services: Air Traffic Control Communication Navigation Airspace management	Information/Feedback Info (e.g. flight plans, schedules, status) Information for use in Air Traffic Control (e.g. ride reports) Finance (via Aviation Trust Fund)
Airports	Industry regulation, certification, and oversight Navigational services Grants in Aid (for safety, capacity, security) [2]	
Business	Revenue	Miscellaneous business services and products
Federal Government	Tool for policy implementation Information (official reports)	Operational budget and laws regarding Aviation Trust Fund [5] FAA appointments (mandate)
Local/Regional/State Government		
Passengers		Revenue (via Aviation Trust Fund) [6]
Society	Wages Mitigation of pollution externalities Information (official reports)	Labor

Figure AP5-05: Stakeholder relationships, Federal Government

Relationship	Federal Government provides (to)	Federal Government receives (from)
Air Transportation Service Providers (ASPs)	Threat of economic regulation Security (e.g. TSA) General services	Taxes (via Treasury) Federal lobby (via lobbying organizations)
Airports	Threat of regulation and oversight (e.g. TSA) General services	Taxes (via Treasury) Critical infrastructure (i.e. national security)
Business	Revenue General services Threat of regulation and oversight	Goods and services Taxes (via Treasury) Federal lobby (via lobbying organizations)
FAA	Operational budget and laws regarding Aviation Trust Fund [5] FAA appointments (mandate)	Tool for policy implementation Information (official reports)
Local/Regional/State Government		
Passengers	Security (e.g. TSA)	Revenue ("September 11th Security Fee" [7])
Society	General services	Taxes (via Treasury) Votes (feedback/mandate) Federal lobby (via lobbying organizations)

Figure AP5-06: Stakeholder relationships, Local/Regional/State Government

Relationship	**Local/Regional/State Government provides (to)**	**Local/Regional/State Government receives (from)**
Air Transportation Service Providers (ASPs)	Tax and other incentives Employment base General services	
Airports	Appointments (mandate) Local infrastructure Financing for construction	Revenue [3] (taxes, other)
Business	General services Revenue Public (local) infrastructure	Taxes Goods and services
FAA		
Federal Government		
Passengers	Travelers assistance and general information Public (local) infrastructure	Revenue (fees, taxes)
Society	General services Public (local) infrastructure	Taxes Votes (feedback/mandate)

Figure AP5-07: Stakeholder relationships, Passengers

Relationship	Passengers provides (to)	Passengers receives (from)
Air Transportation Service Providers (ASPs)	Fares and service fees Information/Feedback (direct/indirect)	Transportation services Promise of future service Amenities and information Contractual obligation
Airports	Revenue (PFC, taxes, fees)	Essential facility services Access to air transportation Amenities, concessions, and parking
Business	Labor Revenue	Wages Goods and services
FAA	Revenue (via Aviation Trust Fund) [6]	
Federal Government	Revenue ("September 11th Security Fee" [7])	Security (e.g. TSA)
Local/Regional/State Government	Revenue (fees, taxes)	Travelers assistance and general information Public (local) infrastructure
Society	<Passengers are a part of society>	<Passengers are a part of society>

Figure AP5-08: Stakeholder relationships, Society

Relationship	**Society provides (to)**	**Society receives (from)**
Air Transportation Service Providers (ASPs)	Labor	Wages Environmental pollution (noise, particulates, GHG)
Airports	Labor	Wages Environmental pollution (noise, particulates, GHG) Information (official reports)
Business	Labor Revenue	Wages Goods and services
FAA	Labor	Wages Mitigation of negative externalities Information (official reports)
Federal Government	Taxes (via Treasury) Votes (feedback/mandate) Federal lobby (via lobbying organizations)	General services
Local/Regional/State Government	Taxes Votes (feedback/mandate)	General services Public (local) infrastructure
Passengers	<Passengers are a part of society>	<Passengers are a part of society>

Stakeholder Relationships, Notes

1: San Francisco Chronicle, 4/1/04; "Airline granted a break UAL given OK to cut payments on SFO bond issue"

2: FAA Budget in Brief, 2008, p. 7

3: San Francisco International Airport Operating Budget FY 2006-7

4: http://www.faa.gov

5:http://www.faa.gov/airports_airtraffic/trust_fund/; 10/20/07

6: AATF funding provided by passenger ticket and jet fuel taxes

7: Aviation and Transportation Security Act passed Nov. 19th 2001; http://www.dot.gov/affairs/dot13001.htm 10/20/07

Bibliography

B. Agle, R. Mitchell, and J. Sonnenfeld, "Who Matters to CEOs? An Investigation of Stakeholder Attributes and Salience, Corporate Performance, and CEO Values," *Academy Of Management Journal*, Vol. 42 No. 5, 1999, pp. 507-525

G. Andreatta and G. Romanin-Jacur, "Aircraft Flow Management Under Congestion," *Transportation Science*, Vol. 21 No. 4, 1987, pp. 249-253

M. Ball, C. Barnhart, G. Nemhauser, and A. Odoni, "Managing Air Traffic and Airline Operations for Schedule Reliability," Operations Research Handbook on Transportation, Laporte and Barnhart eds., in press, 2006

M. Ball, G. Dahl, and T. Vossen, "Matchings in Connection with Ground Delay Program Planning," *Robert H. Smith School Research Paper*, October, 2006

M. Ball, R. Hoffman, D. Lovell, and A. Mukherjee, "Response Mechanisms for Dynamic Air Traffic Flow Management," *6th USA/Europe Air Traffic Management R&D Seminar*, 2005

M. Ball, T. Vossen, and R. Hoffman, "Analysis of Demand Uncertainty Effects in Ground Delay Programs," *4th USA/Europe Air Traffic Management R&D Seminar*, 2001

M. Ball, R. Hoffman, A. Odoni, and R. Rifkin, "A Stochastic Integer Program with Dual Network Structure and its Application to the Ground-Holding Problem," *Operations Research*, Vol. 51 No. 1, 2003, pp. 167-171

C. Barnhart, P. Belobaba, and A. Odoni, "Applications of Operations Research in the Air Transportation Industry," *Transportation Science*, Vol. 37 No. 4, 2003, pp. 368-391

D. Bertsimas and S. Stock-Patterson, "The Air Traffic Flow Management Problem with Enroute Capacities," *Operations Research*, Vol. 46 No. 3, 1998, pp. 406-422

D. Bertsimas and S. Stock-Patterson, "The Traffic Flow Management Rerouting Problem in Air Traffic Control: A Dynamic Network Flow Approach," *Transportation Science*, Vol. 34 No. 3, 2000, pp. 234-255

B. Cameron, "Value Network Modeling," Masters Thesis, Massachusetts Institute of Technology, 2007

K. Chang, K. Howard, R. Oiesen, L. Shisler, M. Tanino, and M. Wambsganss, "Enhancements to the FAA Ground-Delay Program Under Collaborative Decision Making," *Interfaces*, Vol. 31 No. 1, 2001, pp. 57-76

R. E. Freeman, Strategic Management: A Stakeholder Approach, Pitman, Boston, MA, 1984

I. Grossi, "Stakeholder Analysis in the Context of the Lean Enterprise," Masters Thesis, Massachusetts Institute of Technology, 2003

M. Hanowsky, "Efficiency vs. Equity in Dynamic Air Traffic Flow Management," *INFORMS Conference (Seattle, WA)*, November, 2007

R. Hoffman, J. Krozel, and R. Jakobavitz, "Potential Benefits of Fixed-Based Ground Delay Programs to Address Weather Constraints," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, August, 2004

R. Hoffman, J. Krozel, G. Davidson, and D. Kierstead, "Probabilistic Scenario-Based Event Planning for Traffic Flow Management," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, August, 2007

R. Hoffman, A. Mukherjee, and T. Vossen, "Air Traffic Flow Management," in press, September, 2007

R. Hoffman and M. Ball, "A Comparison of Formulations for the Single-Airport Ground-Holding Problem with Banking Constraints," *Operations Research*, Vol. 48 No. 4, 2000, pp. 578-590

P. Kostiuk, E. Gaier, and D. Long, "The Economic Impacts of Air Traffic Congestion," *Air Traffic Control Quarterly*, 1999

B. Kotnyek and O. Richetta, "Equitable Models for the Stochastic Ground-Holding Problem Under Collaborative Decision Making," *Transportation Science*, Vol. 40 No. 2, 2006, pp. 133-146

P. Liu, "Managing Uncertainty in the Single Airport Ground Holding Problem Using Scenario-based and Scenario-free Approaches," Ph.D. Dissertation, University of California, Berkeley, 2007

P. Liu, M. Hansen, and A. Mukherjee, "Scenario-Based Management of Air Traffic Flow: Developing and Using Capacity Scenario Trees," *Transportation Research Board*, Vol. 1951, 2006, pp. 113-121

G. Lulli and A. Odoni, "The European Air Traffic Flow Management Problem," *Transportation Science*, Vol. 41 No. 4, 2007, pp. 431-443

R. Mitchell, B. Agle, and D. Wood, "Toward a Theory of Stakeholder Salience: Defining the Principle of Who and What Really Counts," *Academy Of Management Review*, Vol. 22 No. 4, 1997, pp. 853-886

A. Mostashari, "Stakeholder-Assisted Modeling and Policy Design for Engineering Systems," Ph.D. Dissertation, Massachusetts Institute of Technology, 2005

A. Mukherjee and M. Hansen, "A Dynamic Stochastic Model for the Single Airport Ground Holding Problem," *Transportation Science*, Vol. 41 No. 4, 2007, pp. 444-456

A. Mukherjee, "Dynamic Stochastic Optimization Models for Air Traffic Flow Management," Ph.D. Dissertation, University of California, Berkeley, 2004

A. Odoni, "The Flow Management Problem in Air Traffic Control," Flow Control of Congested Networks, Odoni, Bianco, and Szego eds., Springer-Verlag, Berlin, 1987

O. Richetta and A. Odoni, "Dynamic Solution to the Ground Holding Problem in Air Traffic Control," *Transportation Research: Part A, Policy And Practice*, Vol. 28 No. 3, 1994, pp. 167-185

O. Richetta and A. Odoni, "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control," *Transportation Science*, Vol. 27 No. 3, 1993, pp. 228-238

M. Terrab and A. Odoni, "Strategic Flow Management for Air Traffic Control," *Operations Research*, Vol. 41 No. 1, 1993, pp. 138-152

T. Vossen and M. Ball, "Slot Trading Opportunities in Collaborative Ground Delay Programs," *Transportation Science*, Vol. 40 No. 1, 2006, pp. 29-43

T. Vossen and M. Ball, "Optimization and Mediated Bartering Methods for Ground Delay Programs," *Naval Research Logistics*, Vol. 53 No. 1, 2005, pp. 75-90

T. Vossen, M. Ball, R. Hoffman, and M. Wambsganss, "A General Approach to Equity in Traffic Flow Management and its Application to Mitigating Exemption Bias in Ground Delay Programs," *Air Traffic Control Quarterly*, Vol. 11 No. 4, 2003, pp. 277-292

T. Vossen, "Fair Allocation Concepts in Air Traffic Management," Ph.D. Dissertation, University of Maryland, 2002

P. Vranas, D. Bertsimas, and A. Odoni, "Dynamic Ground-Holding Policies for a Network of Airports," *Transportation Science*, Vol. 28 No. 4, 1994, pp. 275-291

P. Vranas, D. Bertsimas, and A. Odoni, "The Multi-Airport Ground-Holding Problem in Air Traffic Control," *Operations Research*, Vol. 42 No. 2, 1994, pp. 249-262

M. Wambsganss, "Collaborative Decision Making in Air Traffic Management," New Concepts and Methods in Air Traffic Management, Bianco, Dell'Olmo, and Odoni eds., Springer-Verlag, Berlin, Germany, 2001

U.S. Department of Transportation Volpe National Transportation Systems Center, "Enhanced Traffic Management System (ETMS) Functional Description ," 2004

MIT Lincoln Laboratory, "SFO Marine Stratus Forecast System Documentation," 2004