Passenger-Centric Ground Holding: Including Connections in Ground Delay Program Decisions

by

Mallory Jo Soldner

B.S. Industrial and Systems Engineering, Virginia Tech (2007)

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

Master of Science in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2009

© Massachusetts Institute of Technology 2009. All rights reserved.

Author . . . .
Sloan School of Management
May 14, 2009

Certified by . . . .
Amedeo R. Odoni
Professor of Aeronautics and Astronautics and of Civil and
Environmental Engineering
Thesis Supervisor

Accepted by . . . .
Cynthia Barnhart
Co-Director, Operations Research Center

Passenger-Centric Ground Holding: Including Connections in Ground Delay Program Decisions

by

Mallory Jo Soldner

Submitted to the Sloan School of Management
on May 14, 2009, in partial fulfillment of the
requirements for the degree of
Master of Science in Operations Research

Abstract

This research seeks to address potential "passenger-centric" modifications to the way that ground holding delays are allocated in Ground Delay Programs. The allocation of landing slots to arriving flights during time periods when the overall capacity at an airport is reduced due to adverse weather conditions or other circumstances is a well-studied problem in Air Traffic Flow Management, but not from the passenger's perspective. We propose a Passenger-Centric Ground Holding (PCGH) model, which considers both the number of passengers on flights and, notably, when/if they are making connections. In experimental results, PCGH is shown to lead to slot allocations which are significantly different from those in the currently-used first scheduled, first served (FSFS) approach. A systematic analysis is conducted to determine the impact of PCGH on a variety of airport and airline types. Finally, the effects of a maximum-delay-limiting constraint and the convexity of the cost function are investigated.

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

Acknowledgments

Along the way, many people have provided needed lifts and helpful support over the past two years. First and foremost, I would like to thank my advisor, Professor Amedeo Odoni, without whom this thesis would not have been possible. Amedeo was kind in the face of disappointments and was organized and constructive in helping me move forward. I am also grateful to Cindy Barnhart for her research input and encouragement.

I would next like to thank my family and friends. My parents were always a phone call away, as were Brigid and Rachel. I am especially thankful to Lisa, Andrew, and Blair for their support through thick and thin and for all of our adventures. I am also grateful to Phil and Wei for putting up with me and laughing at my jokes and to everyone else who made the ORC an enjoyable place to work. Last, but not least, I would like to acknowledge my BSSC soccer teams, MBFC and SUFC, for all of the good times and needed breaks. My time at MIT may have been a rocky road (with marshmallow surprises always lurking and definitely a few nuts), but at least it had a sweet finish.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In 2007, over 700 million people flew in the United States, generating over 150 billion dollars in airline revenue. The national air transportation system clearly affects the lives of several million Americans every day and is an integral and critical part of the national economy [3]. When severe weather or other circumstances are expected to significantly reduce the runway capacity at an airport, the Federal Aviation Administration (FAA) responds with Air Traffic Flow Management (ATFM) initiatives such as Ground Delay Programs (GDPs) to more effectively allocate runway landing time-slots to flights in dealing with the reduced arrival rate for the airport. GDPs involve assigning ground holding delays at origin airports to flights that are scheduled to land at times when runway capacity at a destination airport is reduced. GDPs are premised on the idea that it is better, in terms of safety and fuel, to hold some of the arriving planes on the ground before they take-off from their origin airports rather than having them circle in the air above the destination airport, unable to land due to the reduced capacity of the arrival runway. Currently, delay is allocated to flights in GDPs, without regard to the number of passengers onboard each flight and to their itineraries (nonstop vs. connecting), in essentially a first scheduled-first served manner. Our research takes the passengers on each flight into consideration and has the potential to lead to GDPs that may be more attractive and effective from the passenger's perspective.

1.1 Outline of the Thesis

Chapter 2 is a review of the main articles and theses in the literature on ground holding and passenger delays. In Chapter 3, a single airport deterministic integer program from Terrab and Odoni in [12] is introduced. Modifications are described to change it into a Passenger-Centric Ground Holding (PCGH) model, where passenger delays are incorporated by considering, for each flight, the number of passengers on board and when/if each of these passengers is making a connection. Additionally, optional constraints are introduced to force a more equitable treatment of aircraft types and airlines.

Chapter 4 contains the results and analysis of our research. A day of landing operations is considered at an airport with deterministic capacity constraints on the number of landings possible per time period. First, a base case is presented which shows that, if passenger costs are taken into consideration, a PCGH allocation may result in solutions with significantly lower costs than those from a first scheduled, first served (FSFS) allocation. The remainder of the chapter systematically explores the impact of the PCGH allocation on the main types of airports and airlines. We examine three major types of airports: a non-hub, a hub with one dominant airline, and a hub with two dominant airlines. We also examine airlines with majority, minority, and equal stakes in an airport; airlines using banks in their scheduling; airlines with different amounts of scheduled time between connections; and airlines operating a fleet consisting of only a single type of aircraft. In addition, the effects of a maximum-delay-limiting constraint and the convexity of the cost function are analyzed. Finally, Chapter 5 summarizes the conclusions of the thesis and describes opportunities for future research.

Chapter 2

Literature Review

A main focus of research in Air Traffic Flow Management (ATFM) is how to respond when inclement weather or other circumstances reduce the number of aircraft that can safely land at an airport. To make decisions of this type, the FAA implements strategic programs, such as Ground Delay Programs (GDPs). The problem of how much ground delay should be assigned to each flight during a GDP is typically referred to as the Ground Holding Problem (GHP).

Over the past 20 years, much research has been conducted on the GHP. Most of the research has dealt with minimizing the sum of the cost of delays to a given set of flights, with the cost to each flight viewed separately and based on factors such as aircraft size and/or other characteristics. Most of the academic literature acknowledges that, if the passengers on each flight were to be taken into account, the allocation of available arrival slots to flights would probably be seriously affected. However, little detailed research has been dedicated to date to a "passenger-centric" investigation of the GHP. To address this area, our research focuses on how to directly integrate the costs of passenger delays, including those due to missed connections, into GDP and GHP decision-making.

In the academic literature, ground holding in ATFM for a single airport was first introduced by Odoni in 1987 in [9]. Since then, extensions of the problem and its modeling have moved in a number of different directions, each adding more details of the complex reality of the National Air Transportation System (NATS). These extensions can be classified with respect to four main attributes: stochasticity, control, scope, and equity.

The first classifier, stochasticity, refers to whether or not the probabilistic nature of airport capacity is recognized and incorporated by the relevant model. Capacities can be deterministic or stochastic. If deterministic, it is assumed that the capacity at the airport under consideration is known. If stochastic, it is assumed that there is a known probability distribution for a number of static or dynamic capacity profiles. In 1993, Terrab and Odoni [12] introduced both a deterministic integer program (IP) and a static-stochastic dynamic program for the single airport ground holding problem. For our research, we are assuming that capacity is deterministically known, and we have based our model on the deterministic IP in [12].

The second classifier, control, refers to how the decision-making responsibilities are allocated between the FAA and the airlines. Decision-making in ground holding can be centralized, with all of the control being held by the FAA, or partially decentralized, in which case the airlines and the FAA share the control of scheduling flights to landing slots in a pre-established system of allocation and mediation known as Collaborative Decision Making (CDM). Centralized control more naturally ties into mathematical modeling – and it is the paradigm we explore in our research. However, decentralized control shared between the FAA and the airlines is more realistic. In [7], after introducing a dynamic-stochastic program with non-linear costs, Hanowsky investigates the trade-off between centralized and decentralized control from the passenger's perspective. For the examples he considered, Hanowsky found that centralized control performed significantly better than what was possible with decentralized control.

The third classifier is scope, which in this setting refers to whether the ground holding decisions are considered at a single airport or throughout the whole network of airports in the NATS. The standard approach in ATFM is to consider only a single airport at a time, and this is what the formulation in [12] does. Bertsimas and Stock Patterson were the first to examine a full network, including air sector capacities in addition to airport capacities, in [4]. Their model built on the multi-airport GHP developed by Vranas, Bertsimas, and Odoni in [13]. Our model will focus on the single airport problem.

The fourth classifier, equity, refers to whether a model explicitly attempts to distribute delay equitably among the stakeholders. The airlines are the stakeholders most prominently taken into account in the research and in actual GDP decisions, as they are constantly vying for shared resources. If a policy in a GDP is perceived as not treating all of the airlines fairly, then the airlines that expect to find themselves at a disadvantage will resist it: even a slight advantage or disadvantage can lead to large economic benefits or costs for an airline. This is an area where CDM has succeeded in building a consensus. Its foundational principle is that landing slots are assigned in a first scheduled, first served (FSFS) manner known as Ration by Schedule (RBS). The airlines consider this method to be fair. Once the airlines have an initial allocation of slots, they can make swaps and cancellations within their set of flights and slots.

The equitable treatment of passengers is a topic much less considered, even though passengers are also important stakeholders in the NATS. In a chapter of [7], Hanowsky considers flight cost functions that are proportional to the number of passengers per flight. However, his research fails to consider what has widely been identified in the literature as the crux of the true delay costs of passengers on flights – missing their connections. Our research considers both the number of passenger on a flight and their connections.

Another interesting approach in considering the costs of passenger delay is presented by Bratu and Barnhart in [6] in 2006. Bratu and Barnhart provide an approach for airline schedule recovery in which the objective is to find the optimal trade-off between airline operating costs and passenger delay costs. Their focus is not on centralized decision-making but on a single airline's specific response to a GDP (and resulting RBS allocation) or to other non-routine disruptions of their daily operations.

Chapter 3

Description of the Problem and Model Formulations

In this chapter, the foundations and development of our research are presented. Specifically, we propose a "passenger-centric" model for the single airport ground holding problem (SAGHP) based on a more accurate assessment of the costs incurred by the airlines and the passengers affected by delays.

The Terrab-Odoni deterministic model, which our new model is based on, is presented in Section 3.1, with the integer programming (IP) formulation in Section 3.1.1 and a discussion of why the IP can be solved as a linear program (LP) in Section 3.1.2. Next Section 3.2 explains how current approaches used in addressing the SAGHP fail to take into account passenger-specific delay information, often resulting in a significant underestimation of the delay costs incurred. The section then introduces a new cost function that more accurately expresses true passenger delay costs in a meaningful but computationally tractable manner. Finally, the Passenger-Centric Ground Holding (PCGH) model is introduced in Section 3.3 with a discussion of additional constraints that can be added to the formulation to enforce a more equitable treatment of aircraft types and airlines.

In both the Terrab-Odoni model and our new PCGH model, we examine the arriving flights at a single airport, airport Z, during a time when the runway capacity is decreased due to weather or some other source of disruption. We then seek to create a feasible schedule of arriving flights, such that there are no delayed flights held in the air, unable to land, above airport Z. This means that all delay is served on the ground before take-off at the various airports from which flights to airport Z are departing.

The model is grounded in three assumptions. We are assuming deterministic knowledge of arrival capacities at airport Z (in terms of the number of flights that can land at any specific time period) and deterministic knowledge of the travel times of the aircraft between each origin and airport Z. We are also assuming that congestion at airport Z is the only cause of delay to incoming flights.

3.1 The Terrab-Odoni Deterministic Model

As noted, the PCGH model proposed in this thesis is a modification of the deterministic model proposed by Terrab and Odoni in [12]. This section of the thesis describes the original integer program (IP) proposed by Terrab and Odoni in 1993. The reader is encouraged to consult the original paper if greater detail is desired.

3.1.1 Model Formulation

The Terrab-Odoni deterministic IP was designed to assign flights to landing slots in the SAGHP. There are N total flights to be scheduled, and I={1,...,N} is the set of these flights, indexed by i. The time interval during which flights from I are originally scheduled to land is subdivided into P time periods of equal length. J={1,...,P+1} is the set of these time periods with the addition of time period P+1. It is assumed that airport Z's arrival capacity during time period P+1 is large enough so that any flights that were not able to land during time periods 1, 2,..., P will be able to land during time period P+1. The set J is indexed by j.

The decision variables, x_{ij}, assign each flight i to land during some time period j, where j must be equal to or later than the time slot when flight i was originally scheduled to land. Once the assignments have been made, the take-off time for any flight i can then be determined. Since we know deterministically in advance the time needed for flight i to travel to airport Z, the take-off time can be calculated by subtracting the flight time from the scheduled landing time.

The following is the IP's formulation:

$$\min \quad \sum_{i=1}^{N}\sum_{j=1}^{P+1} C_{ij}x_{ij} \tag{3.1}$$

$$\text{s.t} \quad \sum_{j=P_i}^{P+1} x_{ij} = 1, \quad \forall i \in \{1, ..., N\} \tag{3.2}$$

$$\sum_{i=1}^{N} x_{ij} \leq K_j, \quad \forall j \in \{1, ..., P\} \tag{3.3}$$

$$x_{ij} \in \{0,1\} \quad \forall i \in \{1, ..., N\} \text{ and } \forall j \in \{1, ..., P+1\} \tag{3.4}$$

where x_{ij} = 1 if flight i is scheduled to land during time period j; 0 otherwise

C_{ij} = flight delay cost of assigning flight i to land during time period j

P_i = time period in which flight i is originally scheduled to land

K_j = arrival capacity of the airport (in no. of flights) during time period j

The objective function, (3.1), states that the objective of the model is to minimize the total cost of the scheduling assignments. The first constraint, (3.2), ensures that every flight eventually lands. In addition, by summing x_{ij} from P_i to $P+1$, a flight i cannot be assigned to land in a time period earlier than when it is originally scheduled to arrive, P_i. Constraint (3.3) ensures that for every time period j, the arrival capacity, K_j, is not exceeded. Lastly, Constraint (3.4), forces the decision variables to be binary.

In the results and analysis of [12], C_{ij} is a flight delay cost. Terrab classified flights into three types: flight with regional jets (RJs), with narrow body jets (NBs), and with wide body jets (WBs). For each flight i, the delay cost over the time horizon

$P_i, ..., P+1$ was based solely on aircraft type. Cost rates were based on rough estimates of actual aircraft costs (i.e. fuel, maintenance, and crew costs) at the time. RJs were the least expensive to delay; WBs were the most; and NBs were in between.

3.1.2 Solving as an LP Results in Integral Solutions

Because their feasible regions are often more complex than those of LPs, IPs generally take much longer to solve than LPs. For this reason, it is important to understand whether the Terrab-Odoni IP can be solved as an LP. As was noted in Terrab's PhD thesis, [11], the constraint matrix for the IP is totally unimodular. Because of this, the problem can be relaxed to an LP by changing (3.4) to $0 \leq x_{ij} \leq 1, \forall i, j$, and it will still yield binary solutions.

Figure 3-1 shows how the IP can can be formulated as a minimum cost network flow problem, i.e., a problem for which the node-arc incidence matrix is totally unimodular (see Theorem 11.12 in [2]) and therefore unimodular. Because the matrix is unimodular and the right-hand-side vector, K_j, is integral, all basic feasible solutions of the LP will be integral (see Theorem 11.11 in [2]).

Further detail on the theory behind total unimodularity and optimization can be found in most optimization textbooks. The reader is referred to [2], Section 11.12 and [5], Section 7.3 for more details.

3.2 Incorporating Passenger-Centric Considerations

In this section, we discuss why the simple metric used currently to account for passenger delays during GDPs fails to capture the impact of missed connections, and we suggest a more accurate delay metric.

3.2.1 Current Deficiencies in Passenger Delay Metrics

The metric typically used to quantify passenger delay in the NATS is "passenger delay-minutes." GHP models that account for passenger delay costs typically com-

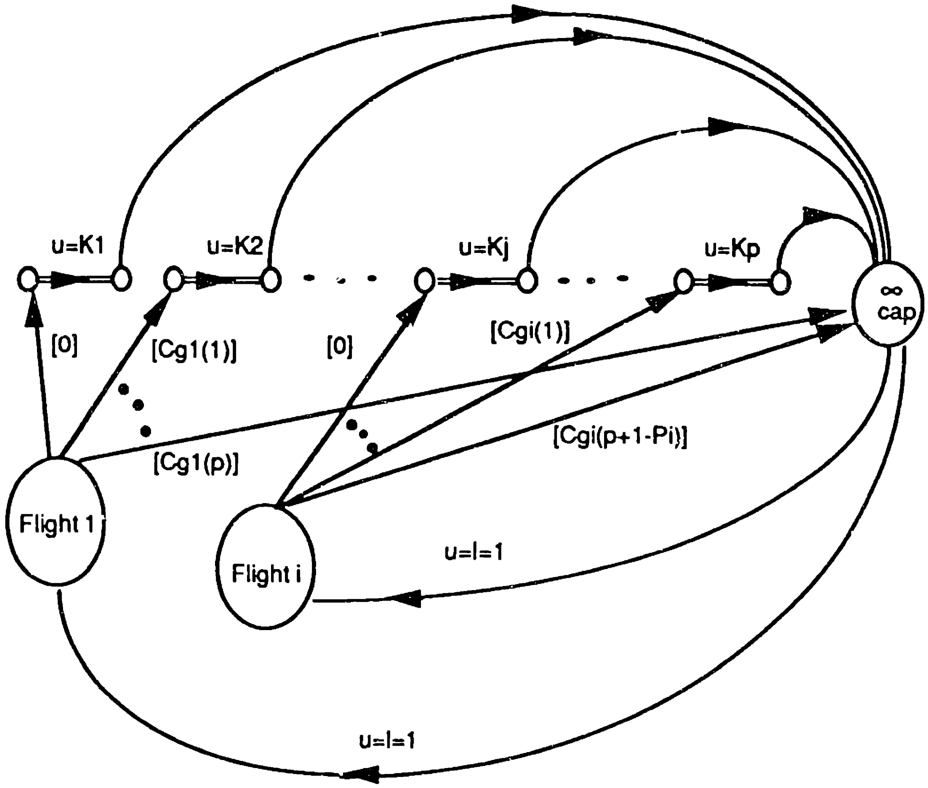

Figure 3-1: Figure 2-1 from [11], which shows how the problem can be expressed as a minimum cost network flow problem. In the figure, costs are in square brackets and upper and lower bounds on capacities are given by u and l, respectively. Cgi(t) is the cost of delaying flight i for t time periods on the ground; it translates to $C_{i,(t+P_i)}$.

pute passenger delay minutes by multiplying the number of passengers onboard a flight by the amount of time the flight is delayed. It is important to note, however, that, although this is an accurate measure for passengers on nonstop flights, the metric falls short for connecting passengers. For a connecting passenger, the measure does not account for the impact of missing a connecting flight. Accordingly, we claim that a more accurate delay metric for passengers is *how late a passenger arrives at his or her destination.*

To understand the difference in the metrics, consider the following examples:

1. Passenger A is a non-connecting passenger, with respect to airport Z. This means that Passenger A's final destination is airport Z, and her trip is a single nonstop leg from some origin airport to airport Z. If her flight is delayed, how late her plane arrives at airport Z is the same as how late she arrives at her final destination.

2. Passenger B is a connecting passenger, again with respect to airport Z. This means that Passenger B will stop at airport Z on the way to his final destination. Passenger B is on a two-leg trip. The first leg of the trip is a flight from his origin airport to airport Z. The second leg is a connecting flight from airport Z to Passenger B's final destination. His first flight is delayed by 30 minutes. However, his original schedule had one hour built-in between flight legs at airport Z, and he makes his connection. Despite one leg of his trip being delayed, he arrives on-time with no delay to his final destination (assuming the final leg is on-time).

3. Passenger B's delay story would have been quite different if the delay on his first flight had caused him to miss his connecting flight. To illustrate this point, consider Passenger C, a connecting passenger with a schedule similar to Passenger B's. Instead of a 30-minute delay, she is delayed for an hour and is not able to make her connection. The next flight to her final destination does not leave for another three hours, so she ends up arriving at her final destination three hours later than scheduled. This delay is different from the one hour of

delay that would have been recorded if one just added up the delays of Passenger C's specific flights. For Passenger C, the passenger delay-minutes metric fails by a wide margin to provide an accurate indication of the delay experienced.

3.2.2 Passenger Delay Cost Function

For passengers who are making connections, the outcome of a delay to their incoming flight is essentially an "all or nothing" proposition. Up to a certain amount of delay, they will still make their connection on an outbound flight and thus incur minimal costs, as was the case for Passenger B. However, when the delay to the incoming flight exceeds a threshold, a missed connection will result, as it did for Passenger C, and the passenger's delay cost incurs a large step of increase. Dealing with this issue adds to the complexity of the ground holding assignment problem: different connecting passengers have different connection times, and considering all the relevant details can become computationally intractable.

The new model proposed in this thesis addresses the issue of delay costs for connecting and non-connecting passengers in a workable way. The idea behind the method is that for a given flight, the aggregate passenger delay cost over the time horizon after its scheduled arrival is the sum of convex functions, one convex function for each passenger on the flight. These convex functions are of two types, one for non-connecting passengers and one for connecting passengers.

Non-Connecting Passengers

The first type of convex function is for non-connecting passengers, such as Passenger A. For these passengers, delay minutes and delay costs increase linearly at a constant rate per minute. These costs are represented by a slightly super-linear function to encourage equity in the assignment of delays among flights with the same passenger and cost profiles.

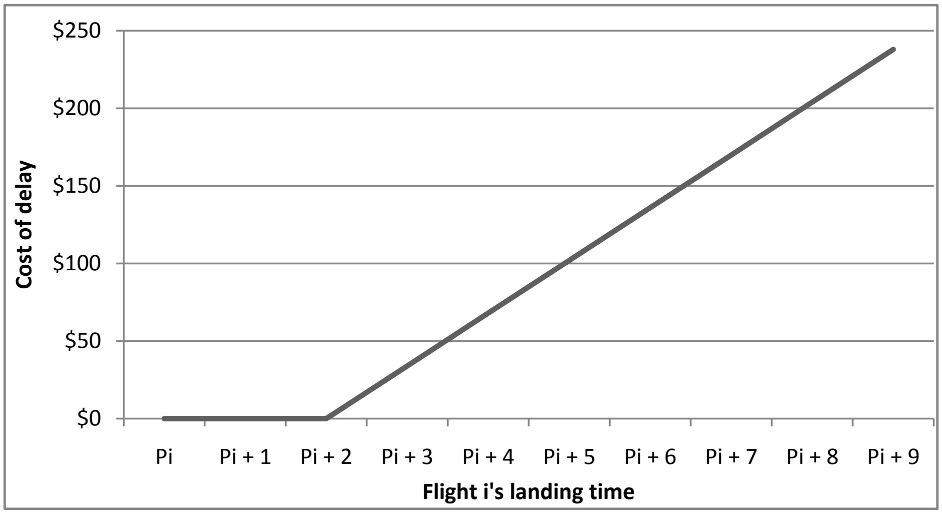

Figure 3-2: Delay cost for a passenger who will miss his/her connecting flight if flight i is delayed three time periods or more.

Connecting Passengers

The other type of convex function is for passengers who are making connections. For these passengers, there is a point in time before which they will not miss their flight and after which they will. The slope of the function up until the breakpoint is zero, and the slope after the breakpoint is significantly higher than that of non-connecting passengers reflecting the immediate effect of the longer delay that accompanies missing one's flight at an airport. Figure 3-2 shows an example of a cost function for a passenger who will miss his or her connecting flight, if flight i is delayed three time periods or more.

Note that in Figure 3-2 and in the other figures presenting cost functions in this thesis, the functions are continuous. However, since time is discretized in the models used in our research, these costs should be thought of as step functions of the actual inputs for the model. Since the size of the steps depends on how long a time period is, displaying them as continuous functions allows for more generality.

It should also be noted that the connecting passenger cost function has flexibility in the steepness of the slope after the breakpoint. Missing a connection is not equally bad for all passengers; the time until the next flight from airport Z to the passenger's final destination can vary widely. For example, consider the difference between missing an international flight that operates only once a day versus missing a regional shuttle that operates every hour. Although this level of detail was not examined in this thesis, the airlines would presumably have information of this type available.

3.3 Passenger-Centric Ground Holding (PCGH) Deterministic Model Formulation

Once more accurate, flight-specific passenger delay cost functions are obtained by adding up the individual cost functions for passengers on the flights, a Passenger-Centric Ground Holding (PCGH) model can be used to determine the ground holding times and landing slot assignments. The formulation of the model is given in this section. The PCGH network flow model is a slight modification of the Terrab-Odoni model presented in Section 3.1.1. The key difference is in the objective function where now there is a passenger delay cost in addition to an aircraft delay cost. Another difference is that the model is presented as a linear program (LP) and not as an IP. This is possible since the constraint matrix is the same as the totally unimodular one in the Terrab-Odoni formulation.

As before, there are N total flights to be scheduled, and I={1,...,N} is the set of these flights, indexed by i. The time interval during which flights from I are originally scheduled to land is subdivided into P time periods of equal length. J={1,...,P+1} is the set of these time periods with the addition of time period P+1. It is assumed that airport Z's arrival capacity during time period P+1 is large enough so that any flights that were not able to land during time periods 1, 2,..., P will be able to land during time period P+1. The set J is indexed by j.

The decision variables, x_{ij}, assign each flight i to land during some time period j, where j must be equal to or later than the time slot when flight i was originally scheduled to land. Again, once the assignments have been made, the take-off time for any flight i can then be determined. Since we know deterministically in advance the time needed for flight i to travel to airport Z, the take-off time can be calculated by subtracting the flight time from the scheduled landing time.

The following is the formulation of the PCGH deterministic model:

$$\begin{aligned}
\min \quad & \sum_{i=1}^{N}\sum_{j=1}^{P+1} C_{ij}x_{ij} + \sum_{i=1}^{N}\sum_{j=1}^{P+1} D_{ij}x_{ij} && (3.5)\\
\text{s.t} \quad & \sum_{j=P_i}^{P+1} x_{ij} = 1, \quad \forall i \in \{1,...,N\} && (3.6)\\
& \sum_{i=1}^{N} x_{ij} \leq K_j, \quad \forall j \in \{1,...,P\} && (3.7)\\
& 0 \leq x_{ij} \leq 1 \quad \forall i \in \{1,...,N\} \text{ and } \forall j \in \{1,...,P+1\} && (3.8)
\end{aligned}$$

where x_{ij} = 1 if flight i is scheduled to land during time period j; 0 otherwise

C_{ij} = flight delay cost of assigning flight to land during time period j

D_{ij} = passenger delay cost of assigning flight i to land during time period j

P_i = time period in which flight i is originally scheduled to land

K_j = arrival capacity of the airport (in no. of flights) during time period j

The objective function, (3.5), states that the objective of the PCGH model is to minimize the total cost of the scheduling assignments. However, unlike in the Terrab-Odoni model, passenger delay costs are now included. As before, the first constraint, (3.6), ensures that every flight eventually lands, and by summing from P_i to $P+1$, a flight i cannot be assigned to land in a time period earlier than when it is originally scheduled to arrive. Constraint (3.7) ensures that for every time period j, the arrival capacity, K_j, is not exceeded.

Lastly, Constraint (3.8), forces the decision variables to fall between zero and one.

Since the constraint matrix is the same in the PCGH formulation as in the Terrab-Odoni model, total unimodularity of the constraint matrix ensures that there exists a binary optimal solution to the LP. If the LP is solved with the network simplex algorithm, a binary optimal solution will be obtained.

3.3.1 Model Extensions

In this subsection, two constraints are introduced that could be added to the PCGH model for added control over the delay allocations. The first constraint limits the maximum delay allowed for any flight or set of flights. The second constraint aims at ensuring an adequately equitable treatment of the airlines by the PCGH model.

Maximum Delay Limitations Constraint

Constraint (3.6) can be modified so that the maximum delay allowable to any flight i is controlled. This modification is shown in (3.9), where M_i is the maximum number of time periods that flight i is allowed to be delayed. M_i can vary by flight.

$$\sum_{j=P_i}^{P_i+M_i} x_{ij} = 1, \quad \forall i \in \{1, ..., N\} \tag{3.9}$$

Constraint (3.9) ensures that no more than M_i time periods of delay are assigned to any flight i. The constraint could be used to give special treatment to a specific flight or group of flights. For example, for flights exempt from ground holding (e.g. international flights), M_i can be set to zero. This ensures that such flights are scheduled to arrive during P_i. Constraint (3.9) could also be used for all flights to ensure that no flight receives more than M_i units of delay. Computational tests of this nature are presented in Section 4.6.

Further, if it is desired to control the maximum delay by arrival time period instead of flight number, a simple calculation for M_i can be performed to determine the model inputs. Let H_j be the maximum number of time periods that any flight scheduled to land at time period j is allowed to be delayed. Let I_{ij} be an indicator of whether flight i is scheduled to land at time j: $I_{ij} = 1$ if $P_i = j$ and 0 otherwise.

(3.10) below can then be used to set M_i. Using M_i from (3.10) in Constraint (3.9) will constrain the maximum number of time periods of delay allowable for any flight scheduled to land during any given time period j.

$$M_i = \sum_{j=1}^{P+1} I_{ij} H_j, \quad \forall i \tag{3.10}$$

It should be noted that using the maximum delay constraint can cause the problem to become infeasible, so M_i needs to be chosen realistically. In addition, using the constraint simply removes arcs (from flight i to all time periods where $j \geq P_i + M_i$) from the network flow model, so the problem remains a minimum cost flow problem that can be solved as an LP.

Equitable Treatment of Airlines Constraint

A major concern for the airlines is whether they are treated in a way that they perceive as equitable during GDPs. Using the PCGH model can result in schedule changes for airlines that are not consistent with a Ration by Schedule allocation, where flights are scheduled to land in essentially the order that they are scheduled. Because of this, it is prudent to address the concern of whether the airlines are treated equitably by a passenger-centric allocation of delays.

The constraint below could be added to the model to allow control of the impact to the airlines. We assume L is a set, indexed by l, of all the airlines operating at least one flight during the time horizon being analyzed. For each airline l, we have A_l, the set of flights i operated by airline l. $N_l = |A_l|$ is the number of flights operated by airline l. R_l is the proportion of delay experienced by airline l in a first scheduled, first served (FSFS) delay allocation of landing slots. Finally, ε is the average number of time periods of deviation from a FSFS allocation allowed per flight. ε is constant for all airlines.

$$\frac{|\sum_{i \in A_l} \sum_{j=P_i}^{P+1} (jx_{ij} - P_i) - R_l \sum_{i=1}^{N} \sum_{i=1}^{P+1} (jx_{ij} - P_i)|}{N_l} \leq \varepsilon, \quad \forall l \in L \tag{3.11}$$

Since airlines consider FSFS scheduling to be fair, the metric that is controlled by Constraint (3.11) is a measure of the deviation from a FSFS allocation. This metric is the difference between the amount of delay that airline l is assigned and the expected amount of delay that would be assigned to airline l in a FSFS allocation, normalized by the number of flights flown by airline l. By using the absolute value of the numerator, the constraint ensures that no airline benefits or is penalized by more than ε.

The disadvantage of using this constraint is that the total unimodularity of the constraint matrix is lost. This means that the model can no longer be solved as an LP and instead needs to be solved as an IP, increasing the solution time.

Chapter 4

Results and Analysis

This chapter presents the example cases, analysis, and results of our research. A day of landing operations is considered at an airport with deterministic capacity constraints on the number of landings possible per time period. Since, in many time periods, demand for landing slots exceeds capacity, there is need for a GDP to assign ground holding delays to flights for safe operations.

Section 4.1 discusses the methodology used in the analysis. In Section 4.2, a base case is presented which shows large savings in cost as a result of using a Passenger-Centric Ground Holding (PCGH) allocation instead of a first scheduled, first served (FSFS) allocation. The remainder of the chapter systematically examines the impact of the PCGH allocation on the main types of airports and airlines. We examine three major types of airports: a non-hub, a hub with one dominant airline, and a hub with two dominant airlines (discussed in Sections 4.3, 4.4, and 4.5, respectively). Throughout these sections, we examine airlines with majority, minority, and equal stakes in the airport; airlines using banks in their schedules; airlines with different amounts of time scheduled between connections; and airlines with uniform fleets (i.e. airlines operating a fleet consisting of only one type of aircraft). The impact of including the maximum delay constraint, (3.9) of Section 3.3.1, in the PCGH model is then considered in Section 4.6. The chapter concludes with a discussion of the dependence of the results on the type of cost function used in PCGH in Section 4.7.

4.1 Methodology

A hypothetical single day's schedule at a congested airport was used in all of the analysis. First the demand profile was created. The hourly demand profile used was the same as that in [12]. From the hourly demand profile, arrival times within each hour and aircraft types were assigned to flights. Details of these assignments can be found in Section 4.1.1. Next passenger and aircraft delay cost functions were assigned to each flight. Details of this can be found in Sections 4.1.3 and 4.1.4. Sections 4.3 - 4.5 describe how airlines were assigned to flights to examine the impact of PCGH on different types of airlines at non-hub and hub airports. Specific details are given in each section, as the assignments varied.

4.1.1 Arrival Schedule Creation

Operations during a day at airport Z were analyzed from 7:00 a.m. to 11:00 p.m., since these are the hours during which most arrival demand occurs at most airports. The 16-hour time horizon was discretized into 96 ten minute time periods. The deterministic hourly demand rates (referring to aircraft arrivals only) were taken from [12]. Since Poisson arrivals are uniformly distributed within a given time interval given a known number of arrivals in that interval, the flights per hour were randomly and uniformly assigned to time periods throughout the hour.

Next, the mix of aircraft types, Type 1, Type 2, and Type 3, were assigned randomly to arrivals with probabilities of 0.4, 0.4, and 0.2, respectively. Type 1 refers to regional jets (RJs). Type 2 refers to narrow-body jets (NBs). Type 3 refers to wide-body jets (WBs). Each of these aircraft types was assumed to have a deterministic number of passengers onboard: RJs with 40 passengers, NBs with 120, and WBs with 240. Figure 4-1 shows the hourly demand profile by aircraft type. It can be seen that within each hour the proportions of Type 1, Type 2, and Type 3 flights were roughly the same (40 percent, 40 percent, and 20 percent), subject to the randomness resulting from the probabilistic assignment.

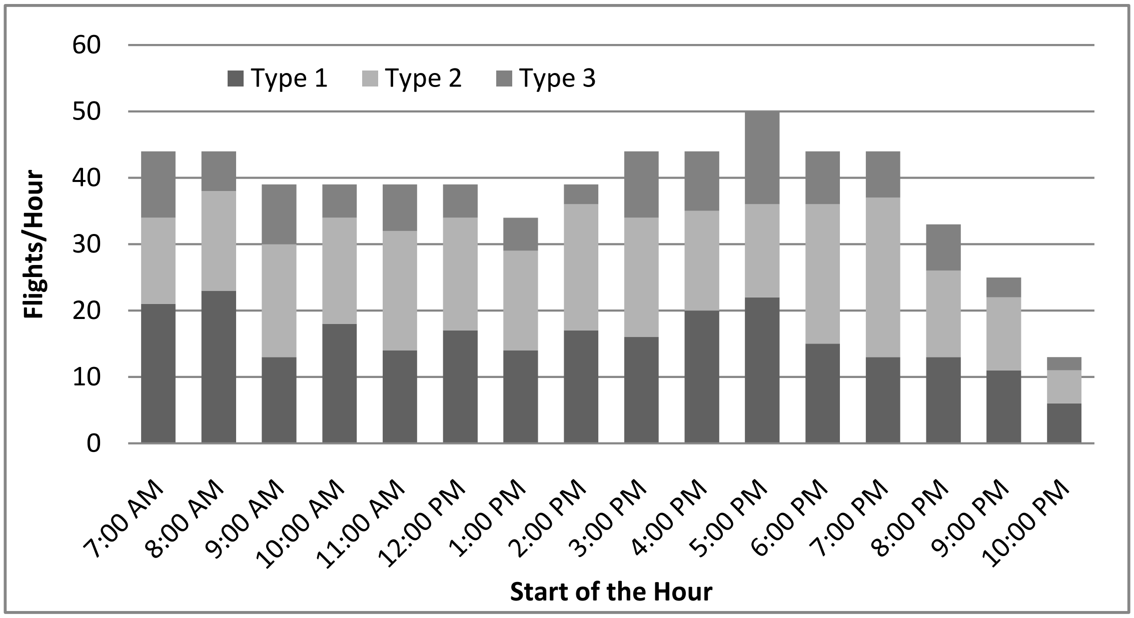

Figure 4-1: Hourly demand profile by aircraft type for arrivals at airport Z.

4.1.2 Capacity

The airport arrival capacity was assumed to be deterministic. Two capacity levels were used in the analysis, one of 45 arrivals per hour and the second of 42 arrivals per hour. These arrival capacities were evenly distributed throughout the time periods in each hour. For 45 arrivals in an hour, K_j, the runway capacity during time period j alternated between seven and eight landings per time period within the hour. For 42 arrivals in an hour, K_j was a constant seven arrivals per time period.

4.1.3 Passenger Delay Costs

Deterministic Treatment of When Passengers Miss Their Connections

For D_{ij}, the passenger delay cost function for flight i, we assumed that there is a known, fixed distribution of the allowable delay before the connecting passengers on an aircraft would miss their connecting flights. We assumed that all flights are full and that 40 percent of the passengers on each flight are making connections. Of the connecting passengers on each flight, it was assumed that, subject to rounding: 5

percent would miss their connections if their incoming flight was delayed 20 minutes; 20 percent would miss if their flight was delayed 30 minutes; 40 percent would miss if their flight was delayed 40 minutes; 25 percent would miss if their flight was delayed 50 minutes; and the final 10 percent would miss if their flight was delayed 60 minutes. Figure 4-2 shows the number of passengers on Type 1, Type 2, and Type 3 aircraft who were assumed to miss their connections per amount of time delayed. For each aircraft type, the same percentage of connecting passengers was assumed to miss their flights at each time period. As the aircraft type increased, the number of passengers missing connections per time period steeply increased. This is due to the differences in total passengers who were assumed to be onboard each aircraft (40 on Type 1, 120 on Type 2, and 240 on Type 3).

The rationale behind the choices of how much delay was allowable for connecting passengers was that passengers would need approximately 30 minutes to travel from the gate at which they arrive to the gate from which they are scheduled to depart on their connecting flight. If the time axis of Figure 4-2 was shifted 30 minutes later, it would represent a realistic distribution of the amount of time scheduled between connecting flights for passengers on most airlines. The distribution in Figure 4-2 is constant in all of the computational tests in this chapter except in Section 4.4.3, where one airline was assumed to have longer scheduled times between connections, shifting the histogram of when passengers were assumed to miss their connections to the right.

The environment in which GDP decisions are made is highly inter-connected, and the assumption of deterministic times until passengers would miss their connecting flights greatly simplified this environment. However, these assumptions make possible the computational tests presented in this thesis by allowing a framework of comparison for many different scenarios. In a real GDP, more accurate data could be available. Whether the airlines would share the information with the FAA for a centrally-controlled GDP is unknown, but the airlines certainly would have available how many passengers were onboard their flights and when/if these passengers were making connections.

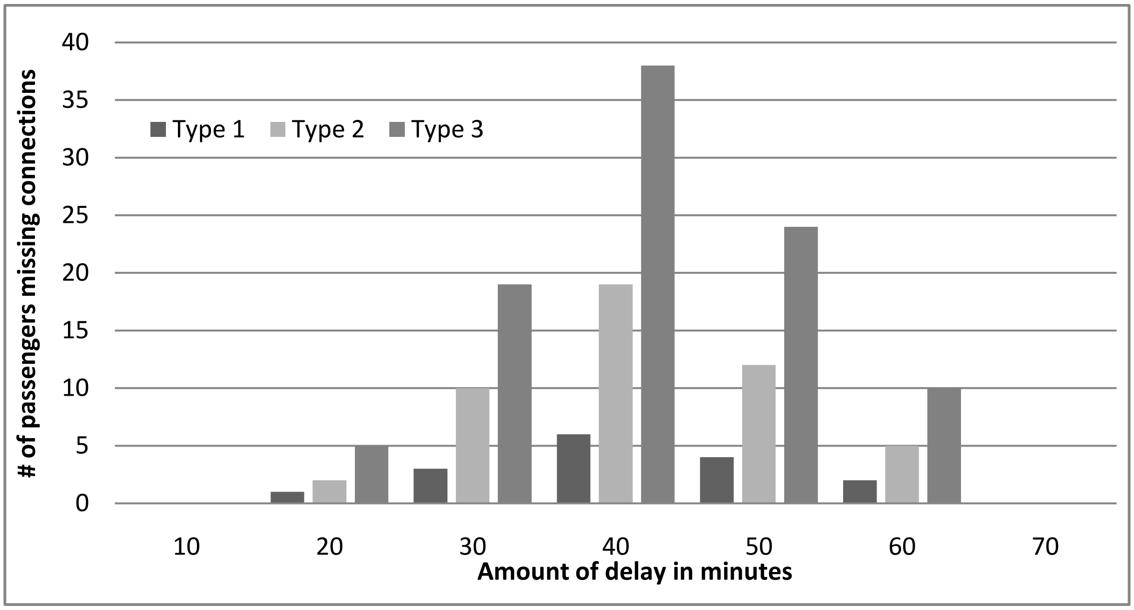

Figure 4-2: Number of passengers for each aircraft type who are assumed to miss their connections per amount of time delayed.

Passenger Delay Cost Functions by Aircraft Type

The total passenger delay cost function for flight i, D_{ij}, is shown in Figure 4-3 for Type 1, Type 2, and Type 3 aircraft. As described in Section 3.2, the cost of passenger delay for a flight is obtained by summing the individual convex delay functions of all of the passengers onboard. These individual cost functions are shown in Figure 4-4. For non-connecting passengers, the slope of the delay cost is five dollars per ten minute time period, which is based on the current standard cost of $30 for delaying an airline passenger for an hour. One source where this delay rate can be obtained is the U.S. Airline Passenger Trip Delay Report [10].

This report estimates that 281.4 million hours of "passenger trip delays" caused an 8.5 billion dollar loss in productivity to the nation's economy in 2006. These numbers reflect that the cost used was $30 per hour of delay (divide 8.5 billion dollars by 281.4 million hours). It should be noted that the 8.5 billion dollar loss in productivity is a conservative estimate since the "passenger trip delays" used in the calculations underestimated true passenger delays since they were based only on single-segment flights and did *not* capture delays due to missed connections. Another useful source

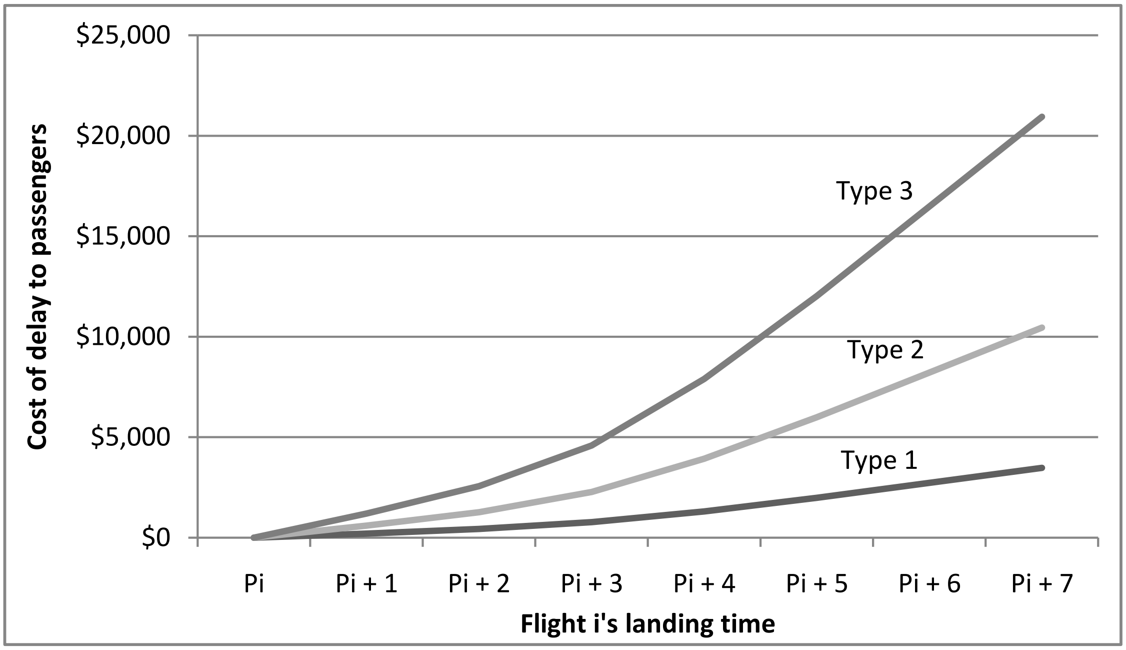

Figure 4-3: Passenger delay costs for some flight i, if i is Type 1, Type 2, or Type 3.

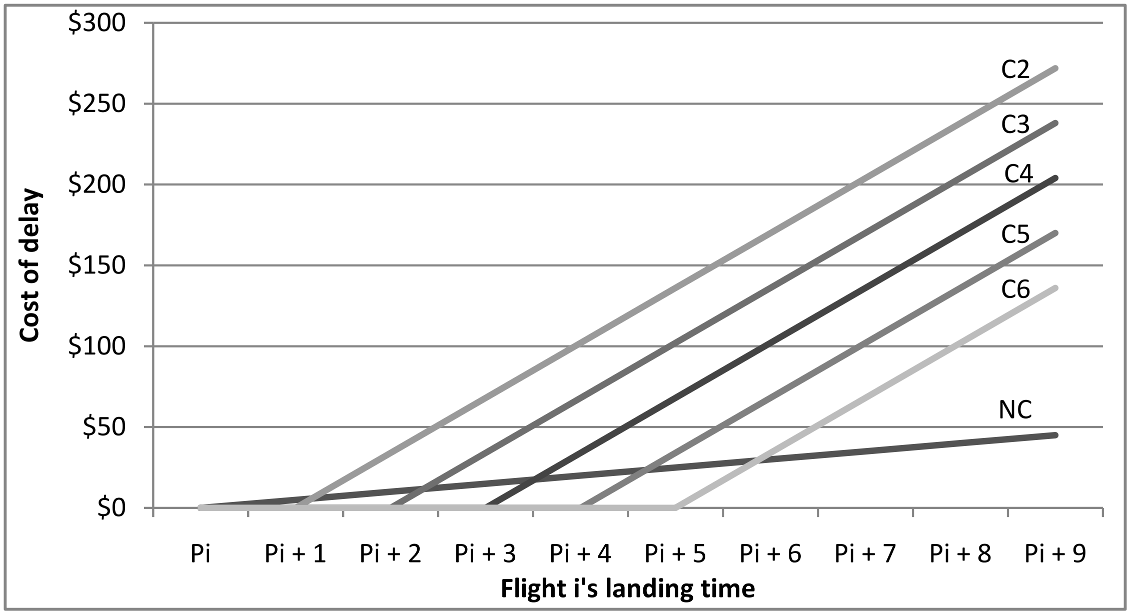

Figure 4-4: Individual cost functions for all passengers on flight i. Passengers are of six types. NC stands for a non-connecting passenger. CX stands for a connecting passenger who will miss his/her connection if flight i is delayed for X time periods or more.

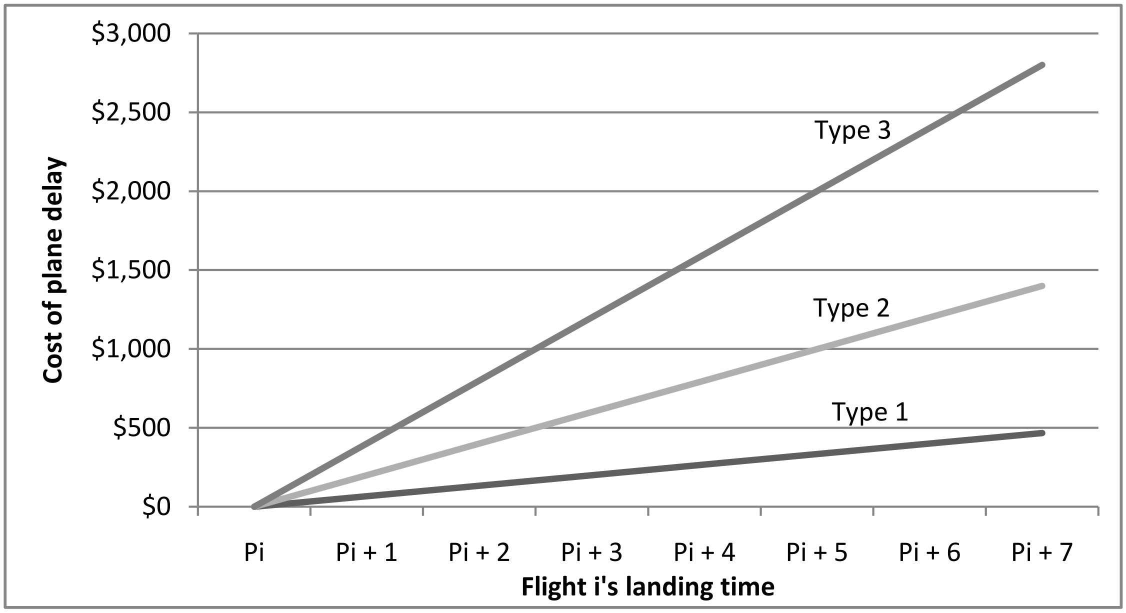

Figure 4-5: Aircraft delay costs for some flight i, if i is Type 1, Type 2, or Type 3.

on the topic is [1].

For connecting passengers, after the flight-missing breakpoint, the slope of the delay cost is 34 dollars per time period. It is intentional that this rate is significantly steeper than the rate for non-connecting passengers. The steep increase in delay cost mimics the immediate step when a passenger misses his or her connection.

4.1.4 Aircraft Delay Costs

The total aircraft delay cost function for flight i, C_{ij}, is shown in Figure 4-5. The hourly delay costs used are the same as in [12]: $400 for an hour of delay to a Type 1 aircraft; $1,200 for an hour of delay to a Type 2 aircraft; and $2,000 for an hour of delay to a Type 3 aircraft. When the computational tests were run, the cost functions were slightly super-linear so that, among aircraft of the same type, it was slightly more expensive to delay one flight by two periods than two flights by one period. This encourages an equitable treatment of individual aircraft in the delay allocations.

			Number of ten minute time periods delayed							
Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
1	253	56.17	110	53	29	32	16	12	1	0
2	251	2.5	236	15	0	0	0	0	0	0
3	110	0	110	0	0	0	0	0	0	0
Total	614	58.67	456	68	29	32	16	12	1	0

Table 4.1: Summary of delay for a PCGH allocation for a runway capacity of 45 landings per hour.

4.2 Base Case

This section presents the delay allocation results for the day of arrival operations without considering how the allocations might impact the airlines (which are not yet assigned to flights). The Passenger-Centric Ground Holding (PCGH) model significantly reduces both passenger and aircraft delay costs when compared with a first-scheduled, first-served (FSFS) model.

4.2.1 PCGH Results

With the passenger and aircraft cost functions and the arrival schedule described in Section 4.1, the PCGH model was used to assign delay allocations to flights for runway capacities of 45 and 42 landings per hour.

Results for the runway capacity of 45 landings per hour are shown in Table 4.1. Type 1 aircraft are delayed more than any other type. Though only 41 percent of flights are Type 1, 96 percent of the total flight-hours of delay are assigned to Type 1 aircraft. In addition, a Type 1 aircraft receives 60 minutes of ground delay, while ten minutes and zero minutes are the longest delays assigned to Type 2 and 3 aircraft, respectively. For this passenger-centric allocation, the total cost of delay for the day is $129,815, $104,348 from passenger delay and $25,467 from aircraft delay.

Results for a runway capacity of 42 landings per hour are shown in Table 4.2. Again, Type 1 aircraft receive most of the delay and longer delays. For this passenger-centric allocation, the total cost of delay for the day is $365,809, $308,476 from passenger delay and $57,333 from aircraft delay.

			Number of ten minute time periods delayed												
Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7	8	9	10	11	12
1	253	110.33	72	53	29	24	13	24	8	7	9	5	6	3	0
2	251	11	185	66	0	0	0	0	0	0	0	0	0	0	0
3	110	0	110	0	0	0	0	0	0	0	0	0	0	0	0
Total	614	121.33	367	119	29	24	13	24	8	7	9	5	6	3	0

Table 4.2: Summary of delay for a PCGH allocation for a runway capacity of 42 landings per hour.

For both of the capacity levels, no Type 3 aircraft are delayed by PCGH. In addition to being the most costly to delay from the perspective of aircraft costs, Type 3 aircraft are also the most costly to delay from the perspective of passengers; there are simply more passengers on WBs than on RJs. This explains why the allocation of delay to Type 3 aircraft is avoided by the PCGH model.

4.2.2 FSFS Results

To obtain an arrival schedule that is FSFS, neither D_{ij} or C_{ij} are used in the objective function. Instead a cost function is used that does not vary by aircraft type and that increases in a slightly super-linear manner. The only factor that impacts the cost is when a flight is originally scheduled to land. As in the case of the aircraft delay costs, discussed in Section 4.1.4, the super-linear increase ensures that it is cheaper to delay two flights by one time period than one flight by two time periods. This creates a FSFS ordering where aircraft that are scheduled to land during earlier time periods are given priority for landing slots.

Results for a runway capacity of 45 landings per hour are shown in Table 4.3. This table should be contrasted with Table 4.1. As expected, the hours of delay are allocated approximately proportionally to the aircraft types (i.e. 41 percent of aircraft are Type 1 and 37 percent of the total flight-hours of delay are allocated to Type 1 aircraft, etc.). Also, the longest that any aircraft is delayed on the ground is 30 minutes, in contrast to the 60 minute longest delay in the PCGH allocation. For the FSFS allocation at capacity 45, the total cost of delay for the day is \$273,081, \$210,014 from passenger delay and \$63,067 from aircraft delay. All of these costs are

			Number of ten minute time periods delayed						
Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6
1	253	21.67	145	86	22	0	0	0	0
2	251	24.5	129	98	23	1	0	0	0
3	110	12.5	49	48	12	1	0	0	0
Total	614	58.67	323	232	57	2	0	0	0

Table 4.3: Summary of delay for a FSFS allocation for a runway capacity of 45 landings per hour.

			Number of ten minute time periods delayed						
Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6
1	253	45.33	109	65	36	37	6	0	0
2	251	51.83	99	65	26	50	11	0	0
3	110	24.17	38	29	16	24	3	0	0
Total	614	121.33	246	159	78	111	20	0	0

Table 4.4: Summary of delay for a FSFS allocation for a runway capacity of 42 landings per hour.

higher than the corresponding costs in the PCGH allocation.

Results for a runway capacity of 42 landings per hour are shown in Table 4.4, which should be contrasted with Table 4.2. As is the case for the FSFS allocation, when the runway capacity is 45 landings per hour, delay is proportional in amount and length by aircraft type. The total cost of delay for the day is \$632,305, \$503,638 from passenger delay and \$128,667 from aircraft delay. Again, these costs are higher than the corresponding costs in the PCGH allocation.

4.2.3 PCGH vs. FSFS

There are two main points of comparison between PCGH and FSFS allocations. The first relates to the overall costs of delays. PCGH reduces delay costs by a wide margin when compared to FSFS. The second point of comparison relates to how the delay allocations impact the different aircraft types. The same total flight-hours of delay are allocated differently among the aircraft types in PCGH than in FSFS.

There are significant savings in passenger, aircraft, and total (the sum of passenger

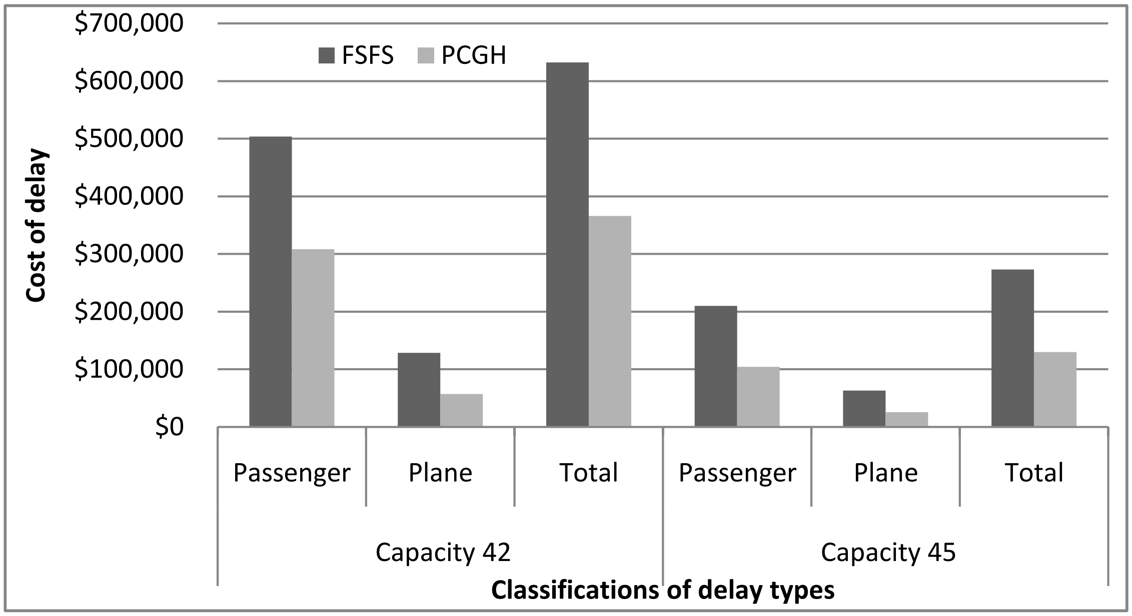

Figure 4-6: A comparison of the delay costs between the FSFS and the PCGH allocations.

and aircraft) delay costs in the PCGH allocations when compared with the FSFS allocations at both capacity levels. Figure 4-6 summarizes these costs savings and provides a visual comparison of the cost breakdowns. In all categories, PCGH results in savings of at least 38 percent. The savings are over 50 percent for aircraft costs at capacity 42 and for all costs at capacity 45. These are significant savings, especially considering that the total flight-hours of delay do not change and that *all of the savings come from reallocating the delay.* It is also seen that passenger delay costs make up a majority of the total delay costs, although this is a direct result of the cost functions used (compare Figures 4-3 and 4-5).

In both FSFS and PCGH allocations, the total flight-hours of delay are equal at the same capacity levels (58.67 hours for capacity 45 and 121.33 hours for capacity 42). The difference between the allocations is in which aircraft types are delayed and for how long. In the FSFS allocation, delay is distributed among Type 1, 2, and 3 flights in proportion to the total number of these aircraft types in the schedule. In the PCGH allocation, the delay is primarily served by Type 1 aircraft and secondarily by Type 2 aircraft.

As a whole and by design, PCGH delays fewer flights than does FSFS. While most Type 2 and 3 aircraft receive shorter delays in PCGH than they would in FSFS, many Type 1 aircraft receive longer delays, where the degree of how much longer is dependent on congestion level. The less congestion, the closer the PCGH allocations are to the FSFS ones, in the sense that the longest delays assigned in PCGH are closer to the longest delays assigned in FSFS. However, at lower congestion levels, the delay is served almost exclusively by Type 1 aircraft, resulting in higher percentages of savings, at the expense of more imbalances in the types of aircraft being delayed. On the other hand, the higher the congestion level, the further the PCGH allocations are from the FSFS ones, in the sense that the longest delays assigned in PCGH are further from the longest delays assigned in FSFS. However, at higher congestion levels, delay is more likely to be assigned to Type 2 and 3 aircraft in PCGH, resulting in a lower percentages of savings than if most of the delay is able to be served by Type 1 aircraft.

4.3 Non-Hub Airport

Once the base case was analyzed, it was desired to judge the potential impact of PCGH on the airlines. This section and the next two systematically consider the impact of PCGH allocations on the main classifications of airports and airlines. In this section, we present the first airport type, a non-hub. The airline types represented are those with equal stakes in an airport and not using bank scheduling.

We assume that there are five airlines operating at airport Z, each operating approximately 20 percent of the arriving flights. For consistency with the results in Section 4.2, the assumption remains that 40 percent of the passengers on each flight are making connections.

4.3.1 The Assignment of Airlines to Flights

Given the fixed schedule of arriving flights and aircraft types (the same schedule described in Section 4.1 and analyzed in Section 4.2), Airlines A, B, C, D, and E were randomly assigned to flights. Each airline had a 20 percent chance of being assigned

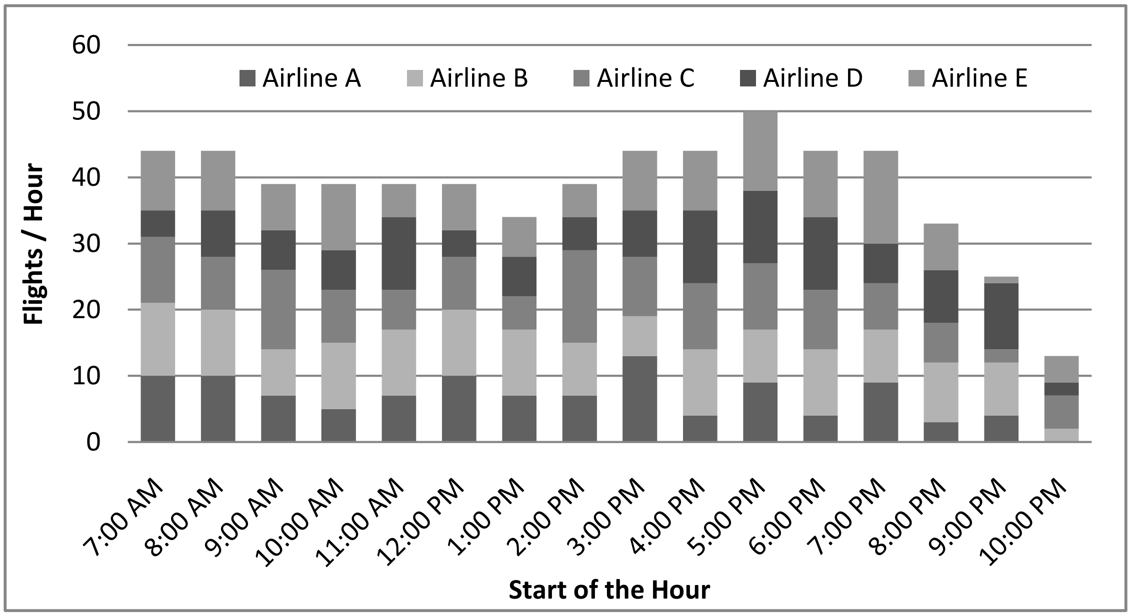

Figure 4-7: Hourly demand profile by airline for arrivals at a non-hub airport.

to each flight i. This random assignment maintained approximately a 40 percent, 40 percent, and 20 percent distribution of Type 1, Type 2, and Type 3 aircraft, respectively, for each of the airlines. Figure 4-7 shows the hourly demand profile by airline. It can be seen that within each hour the proportions of flights to each airline were the same, subject to the randomness resulting from the probabilistic assignment. Assigning airlines to flights in this manner ensured that the airlines were not using banks in their scheduling, a trait desired in this computational test.

A bank, or wave, is when several arrival (or departure) times of arriving (or departing) flights for an airline are scheduled near one another. If an airline is using banks in their scheduling, there will usually be a wave of arriving flights followed by a wave of departing flights. In this way, the airline is better able to manage the flows of passengers from different arriving flights who are making connections to a variety of departing flights. Banks are primarily found at airports which are hubs for an airline, so it is justifiable that the airlines considered in this non-hub airport section are assumed not to use them.

4.3.2 Results

After the airline assignments were made, the models were run to create PCGH and FSFS allocations.

PCGH Allocation

Tables 4.5 and 4.6 contain summaries of the delay statistics by airline for the PCGH allocations at runway capacities of 45 and 42 landings per hour, respectively. The delay allocations are fairly similar among the airlines, especially at capacity 45. This result is expected since the airlines have nearly equal stakes in the airport. It also makes sense that the allocations of delay are more even at the higher capacity level. The less congestion, the less the allocations are subject to inequity due to certain airlines having been randomly assigned to more flights during the more congested times of day.

However, despite the overall similarities in amounts of delay allocated, there are some distinctions. In both cases, Airline E has the most total flight-hours of delay, despite the fact that they do not operate the most flights. On the other hand, Airline B, the airline operating the most flights, is tied for first place for the least number of flight-hours of delay among the airlines at capacity 45. However at capacity 42, Airline B is in second to last place in this ranking. So, how well an airline fares relative to others in the PCGH allocation is not constant across all capacity levels.

FSFS Allocation

Tables 4.7 and 4.8 contain summaries of the delay statistics by airline for the FSFS allocations at runway capacities of 45 and 42 landings per hour, respectively. The delay allocations are not as uniform across the airlines in the FSFS allocations as they are in the PCGH allocations. As in the PCGH allocation, Airline E has the most flight-hours of delay at both capacities, even though the airline does not operate the most flights. This suggests that Airline E's higher delay level is a result of having been randomly assigned to more flights during some of the most congested time periods of

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	43	11.17	13	10	10	5	3	2	0	0
	2	48	0.33	46	2	0	0	0	0	0	0
	3	18	0.00	18	0	0	0	0	0	0	0
Total		109	11.50	77	12	10	5	3	2	0	0
B	1	54	10.67	28	9	4	8	2	3	0	0
	2	62	0.50	59	3	0	0	0	0	0	0
	3	21	0.00	21	0	0	0	0	0	0	0
Total		137	11.17	108	12	4	8	2	3	0	0
C	1	50	10.50	21	13	4	8	3	0	1	0
	2	50	0.67	46	4	0	0	0	0	0	0
	3	29	0.00	29	0	0	0	0	0	0	0
Total		129	11.17	96	17	4	8	3	0	1	0
D	1	54	11.00	27	11	5	4	2	5	0	0
	2	39	0.33	37	2	0	0	0	0	0	0
	3	22	0.00	22	0	0	0	0	0	0	0
Total		115	11.33	86	13	5	4	2	5	0	0
E	1	52	12.83	21	10	6	7	6	2	0	0
	2	52	0.67	48	4	0	0	0	0	0	0
	3	20	0.00	20	0	0	0	0	0	0	0
Total		124	13.50	89	14	6	7	6	2	0	0

Table 4.5: Summary of delay for a PCGH allocation at a non-hub airport with a runway capacity of 45 landings per hour.

				Number of 10 minute time periods delayed												
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7	8	9	10	11	12
A	1	43	20.17	10	8	9	3	2	2	2	3	2	1	1	0	0
	2	48	1.67	38	10	0	0	0	0	0	0	0	0	0	0	0
	3	18	0.00	18	0	0	0	0	0	0	0	0	0	0	0	0
Total		109	21.83	66	18	9	3	2	2	2	3	2	1	1	0	0
B	1	54	21.83	19	11	5	2	4	6	3	0	0	1	2	1	0
	2	62	2.83	45	17	0	0	0	0	0	0	0	0	0	0	0
	3	21	0.00	21	0	0	0	0	0	0	0	0	0	0	0	0
Total		137	24.67	85	28	5	2	4	6	3	0	0	1	2	1	0
C	1	50	20.00	13	14	4	7	2	5	0	0	3	1	0	1	0
	2	50	2.33	36	14	0	0	0	0	0	0	0	0	0	0	0
	3	29	0.00	29	0	0	0	0	0	0	0	0	0	0	0	0
Total		129	22.33	78	28	4	7	2	5	0	0	3	1	0	1	0
D	1	54	21.33	18	11	7	5	1	6	1	0	1	1	2	1	0
	2	39	1.50	30	9	0	0	0	0	0	0	0	0	0	0	0
	3	22	0.00	22	0	0	0	0	0	0	0	0	0	0	0	0
Total		115	22.83	70	20	7	5	1	6	1	0	1	1	2	1	0
E	1	52	27.00	12	9	4	7	4	5	2	4	3	1	1	0	0
	2	52	2.67	36	16	0	0	0	0	0	0	0	0	0	0	0
	3	20	0.00	20	0	0	0	0	0	0	0	0	0	0	0	0
Total		124	29.67	68	25	4	7	4	5	2	4	3	1	1	0	0

Table 4.6: Summary of delay for a PCGH allocation at a non-hub airport with a runway capacity of 42 landings per hour.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	43	4.17	23	15	5	0	0	0	0	0
	2	48	3.33	31	14	3	0	0	0	0	0
	3	18	2.33	5	12	1	0	0	0	0	0
Total		109	9.83	59	41	9	0	0	0	0	0
B	1	54	4.33	31	20	3	0	0	0	0	0
	2	62	5.17	35	23	4	0	0	0	0	0
	3	21	1.00	16	4	1	0	0	0	0	0
Total		137	10.50	82	47	8	0	0	0	0	0
C	1	50	4.00	30	16	4	0	0	0	0	0
	2	50	5.17	26	17	7	0	0	0	0	0
	3	29	3.50	11	15	3	0	0	0	0	0
Total		129	12.67	67	48	14	0	0	0	0	0
D	1	54	4.17	35	13	6	0	0	0	0	0
	2	39	4.50	15	21	3	0	0	0	0	0
	3	22	2.67	9	10	3	0	0	0	0	0
Total		115	11.33	59	44	12	0	0	0	0	0
E	1	52	5.00	26	22	4	0	0	0	0	0
	2	52	6.33	22	23	6	1	0	0	0	0
	3	20	3.00	8	7	4	1	0	0	0	0
Total		124	14.33	56	52	14	2	0	0	0	0

Table 4.7: Summary of delay for a FSFS allocation at a non-hub airport with a runway capacity of 45 landings per hour.

the day, which can be observed at the hourly level in Figure 4-7.

PCGH vs. FSFS

Figure 4-8 shows a comparison of the delay costs associated with a PCGH vs. a FSFS allocation at a non-hub airport with a runway capacity of 45 landings per hour. As seen in the delay tables, the PCGH allocation has more uniform costs across the airlines than the FSFS one. By itself, this could be considered a good characteristic of PCGH, indicating a more equitable treatment of the airlines, since both the ratios of delay costs and of flights operated are closer to one another for each of the airlines in PCGH than in FSFS. However, the consequence of this trait is that the percentage cost savings for each airline from using PCGH instead of FSFS are not uniform, and therefore, not as equitable. These improvements range from a 40 percent reduction in total costs for Airline B to a 62 percent improvement for Airline C.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	43	7.83	18	14	2	7	2	0	0	0
	2	48	8.17	23	13	2	8	2	0	0	0
	3	18	3.67	5	7	3	3	0	0	0	0
Total		109	19.67	46	34	7	18	4	0	0	0
B	1	54	8.50	26	12	9	7	0	0	0	0
	2	62	13.33	24	13	10	13	2	0	0	0
	3	21	2.83	12	4	2	3	0	0	0	0
Total		137	24.67	62	29	21	23	2	0	0	0
C	1	50	8.33	21	15	8	5	1	0	0	0
	2	50	9.17	23	13	4	6	4	0	0	0
	3	29	6.83	8	9	5	6	1	0	0	0
Total		129	24.33	52	37	17	17	6	0	0	0
D	1	54	9.17	26	9	12	6	1	0	0	0
	2	39	9.67	9	15	4	9	2	0	0	0
	3	22	5.83	7	5	1	8	1	0	0	0
Total		115	24.67	42	29	17	23	4	0	0	0
E	1	52	11.50	18	15	5	12	2	0	0	0
	2	52	11.50	20	11	6	14	1	0	0	0
	3	20	5.00	6	4	5	4	1	0	0	0
Total		124	28.00	44	30	16	30	4	0	0	0

Table 4.8: Summary of delay for a FSFS allocation at a non-hub airport with a runway capacity of 42 landings per hour.

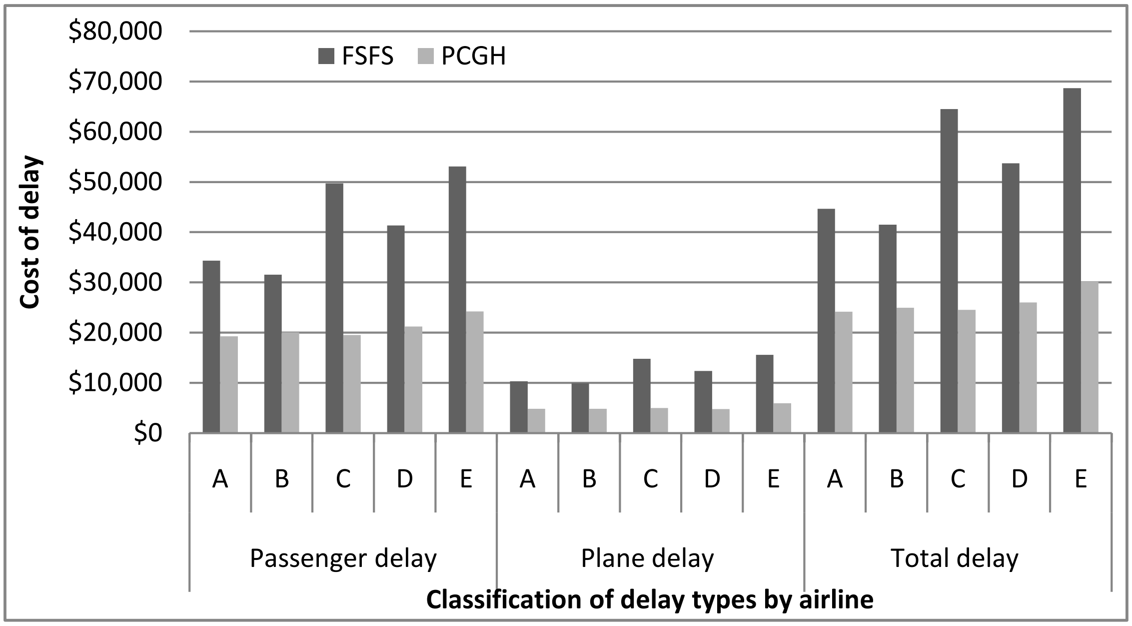

Figure 4-8: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a non-hub airport with a runway capacity of 45 landings per hour.

Figure 4-9 contains a comparison of the delay costs associated with a PCGH vs. a FSFS allocation at a non-hub airport with a runway capacity of 42 landings per hour. Delay costs from the PCGH allocation are not as uniform across the airlines at capacity 42 as they are at capacity 45, as can be seen by contrasting Figure 4-9 with 4-8. Even so, the differences in percentages of delay cost improvements across the airlines are still high, ranging from a 33 percent improvement in total cost for Airline A to a 53 percent improvement for Airline C.

What is generalizable from this section is that PCGH allocations treat equitably, in terms of proportions of delay, airlines with similar schedules, fleets, percentages of connecting passengers, and times between connecting flights. However, this is dependent on how many aircraft of each type each airline operates during the peaks and valleys in the airport demand. FSFS allocations are also dependent on how many aircraft each airline operates during the peaks and valleys in demand, but not on the aircraft types. The difference in aircraft-type dependence between PCGH and FSFS leads to large variations in percentage improvements in costs among the airlines, an

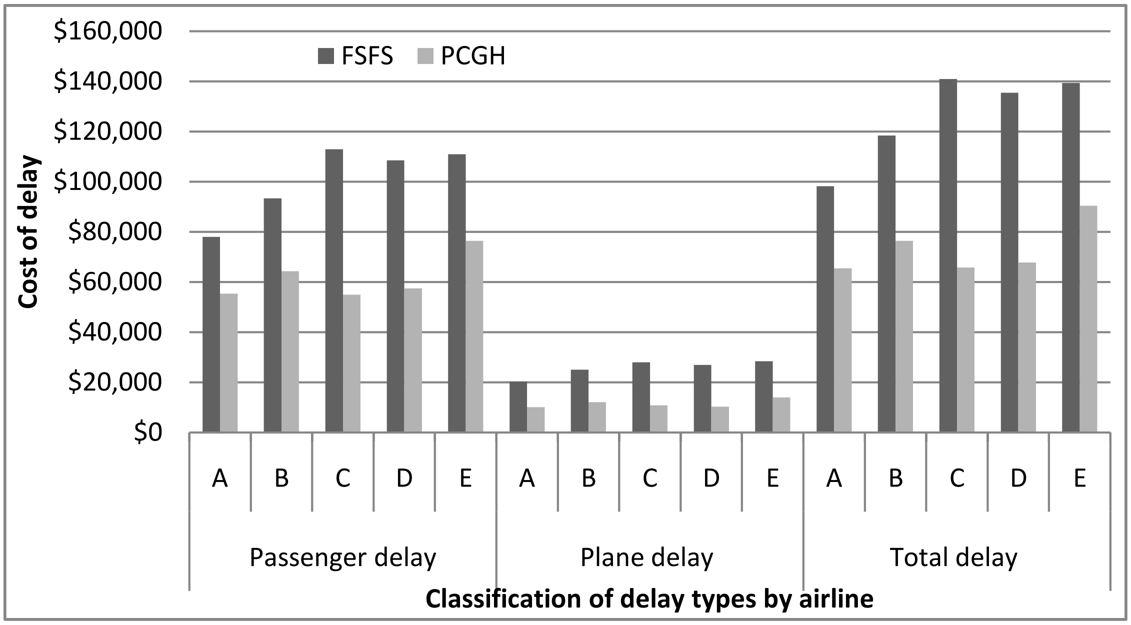

Figure 4-9: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a non-hub airport with a runway capacity of 42 landings per hour.

unattractive quality from the perspective of the airlines.

4.4 Hub Airport with One Dominant Airline

Results for the second of the three airport types are presented next. In this section, we focus on a hub airport with one dominant airline. The types of airlines considered are ones with majority and minority stakes in the airport and ones with longer amounts of time scheduled between connections.

The assignment of airlines to flights is described in Section 4.4.1. Airline A is assumed to operate 80 percent of the flights, and Airline B is assumed to operate 20 percent of the flights in four, one-hour banks. Following the airline assignment subsection, results are presented in two subsections which contrast how the PCGH allocations are affected by how much time the airlines have scheduled between flights for connecting passengers.

Section 4.4.2 focuses on the first case where both airlines have the same distribu-

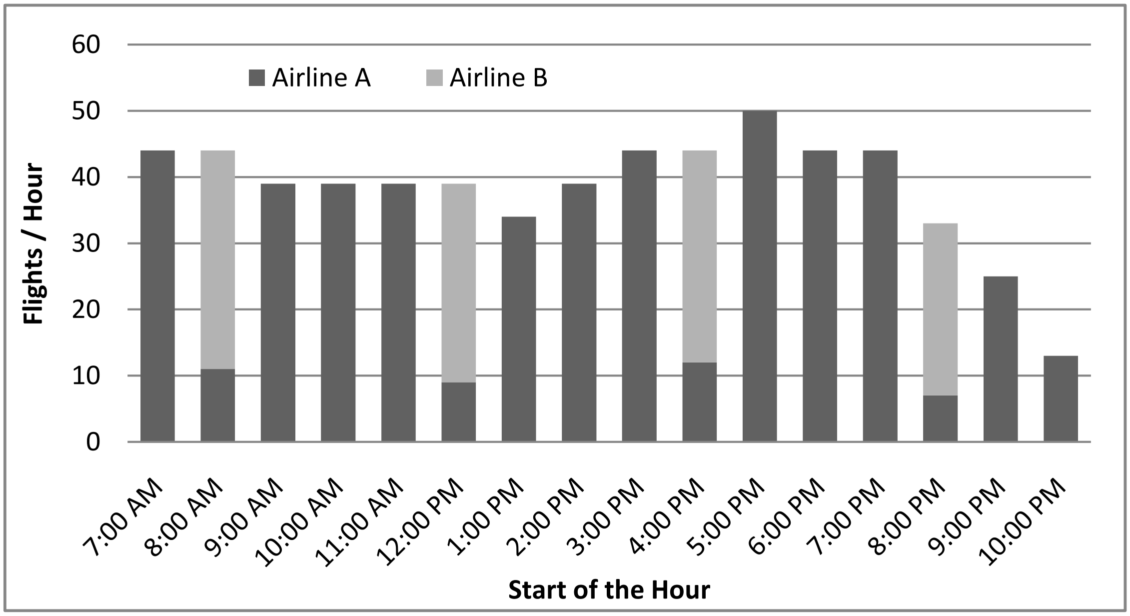

Figure 4-10: Hourly demand profile by airline for arrivals at a hub airport with one dominant airline.

tions of when passengers are assumed to miss their connections. Recall from Section 4.1.3 that a deterministic number of connecting passengers for each flight are assumed to miss their connections after different amounts of allowable delay, as shown in Figure 4-2. Section 4.4.3 focuses on the second case where Airline B is assumed to have longer times built into passengers' itineraries between banks of arrivals and departures, shifting the distribution in Figure 4-2 to the right. The assumption throughout remains that 40 percent of the passengers on each flight are making connections, which is reasonable since both airlines are assumed to use the airport as a hub.

4.4.1 The Assignment of Airlines to Flights

To assign the airlines to the schedule of arriving flights in a way that created a reasonable, hypothetical schedule for the hub airport, first assignments to Airline B were made, leaving the remaining flights for Airline A. For Airline B, four one-hour banks were selected, 8:00 a.m.-8:59 a.m., 12:00 p.m.-12:59 p.m., 4:00 p.m.-4:59 p.m., and 8:00 p.m.-8:59 p.m.. From the pre-determined flight arrival schedule, there were 160 flights scheduled during these four hours. To get a total of approximately 20 percent of all the 614 total flights assigned to Airline B, the 160 flights in the banks were randomly assigned to Airline B, each with probability 0.77 ($0.2 \cdot 614 = 0.77 \cdot 160$). All of the remaining flights were then assigned to Airline A. Figure 4-10 shows the hourly demand profile by airline; Airline B's banks are easily seen. Since airline assignments did not depend on aircraft type, the distribution of aircraft types (40, 40, and 20 percent for Types 1, 2, and 3, respectively) remained about the same for both airlines.

4.4.2 Results for Identical Times Between Connections

Results are first presented for the case when both airlines have the same distributions of when connecting passengers are assumed to miss their flights. As in Sections 4.2 and 4.3, the number of passengers for each aircraft type who are assumed to miss their connections per amount of time delayed follows the distribution shown in Figure 4-2. The models were run to create PCGH and FSFS allocations for this case.

PCGH Allocation

Tables 4.9 and 4.10 contain summaries of the delay statistics by airline for the PCGH allocations at runway capacities of 45 and 42 landings per hour, respectively. At both capacity levels, the proportions of flight-hour delays for the airlines are close to the proportions of total flights flown by each (i.e. Airline B operates approximately 20 percent of the flights and is allocated just under 20 percent of the flight-hours of delay).

FSFS Allocation

Tables 4.11 and 4.12 contain summaries of the delay statistics by airline for the FSFS allocation at runway capacities of 45 and 42 landings per hour, respectively. As in the PCGH allocation, the proportions of flight-hour delays for both airlines are close to the proportions of total flights operated by each.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	199	45.83	92	38	13	27	16	12	1	0
	2	202	2.50	187	15	0	0	0	0	0	0
	3	92	0.00	92	0	0	0	0	0	0	0
Total		493	48.33	371	53	13	27	16	12	1	0
B	1	54	10.33	18	15	16	5	0	0	0	0
	2	49	0.00	49	0	0	0	0	0	0	0
	3	18	0.00	18	0	0	0	0	0	0	0
Total		121	10.33	85	15	16	5	0	0	0	0

Table 4.9: Summary of delay for a PCGH allocation at a hub airport with one dominant airline with a runway capacity of 45 landings per hour, for the case when both airlines have the same distributions of times between connecting flights.

				Number of 10 minute time periods delayed												
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7	8	9	10	11	12
A	1	199	89.67	62	43	19	12	8	20	5	7	9	5	6	3	0
	2	202	9.17	147	55	0	0	0	0	0	0	0	0	0	0	0
	3	92	0.00	92	0	0	0	0	0	0	0	0	0	0	0	0
Total		493	98.83	301	98	19	12	8	20	5	7	9	5	6	3	0
B	1	54	20.67	10	10	10	12	5	4	3	0	0	0	0	0	0
	2	49	1.83	38	11	0	0	0	0	0	0	0	0	0	0	0
	3	18	0.00	18	0	0	0	0	0	0	0	0	0	0	0	0
Total		121	22.50	66	21	10	12	5	4	3	0	0	0	0	0	0

Table 4.10: Summary of delay for a PCGH allocation at a hub airport with one dominant airline with a runway capacity of 42 landings per hour, for the case when both airlines have the same distributions of times between connecting flights.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	199	17.00	118	60	21	0	0	0	0	0
	2	202	19.67	109	69	23	1	0	0	0	0
	3	92	10.50	43	36	12	1	0	0	0	0
Total		493	47.17	270	165	56	2	0	0	0	0
B	1	54	4.67	27	26	1	0	0	0	0	0
	2	49	4.83	20	29	0	0	0	0	0	0
	3	18	2.00	6	12	0	0	0	0	0	0
Total		121	11.50	53	67	1	0	0	0	0	0

Table 4.11: Summary of delay for a FSFS allocation at a hub airport with one dominant airline with a runway capacity of 45 landings per hour.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	199	35.00	94	43	24	33	5	0	0	0
	2	202	41.50	86	47	16	42	11	0	0	0
	3	92	20.00	33	25	10	21	3	0	0	0
Total		493	96.50	213	115	50	96	19	0	0	0
B	1	54	10.33	15	22	12	4	1	0	0	0
	2	49	10.33	13	18	10	8	0	0	0	0
	3	18	4.17	5	4	6	3	0	0	0	0
Total		121	24.83	33	44	28	15	1	0	0	0

Table 4.12: Summary of delay for a FSFS allocation at a hub airport with one dominant airline with a runway capacity of 42 landings per hour.

PCGH vs. FSFS

Figures 4-11 and 4-12 show comparisons of the delay costs associated with a PCGH vs. a FSFS allocation at a hub airport with one dominant airline with runway capacities of 45 and 42 landings per hour, respectively. At both capacity levels, Airline B saves about ten percent more than Airline A. From Figure 4-10, it can be seen that Airline B has only one bank during the hours of highest demand (3:00 p.m. - 7:00 p.m.), and this bank is before the highest peak in demand. This explains why the longest delays allocated to Airline B's flights are shorter than those to Airline A's flights, as can be seen in the right-hand sides of Tables 4.9 - 4.12. While these shorter lengths of delays for Airline B are present in both the PCGH and FSFS allocations, they are more pronounced in the PCGH allocations. The shorter delays to Airline B's flights translate into fewer of the airline's passengers missing their flights and explain why Airline B has higher percentages of savings than Airline A does in Figures 4-11 and 4-12.

4.4.3 Results for Differing Times Between Connections

Second, results are presented for the case when Airline B has longer amounts of time between arriving and departing banks of flights. Unlike in the last subsection, the passengers on Airline B's flights do not follow the distribution in Figure 4-2 of the number of passengers for each aircraft type who are assumed to miss their

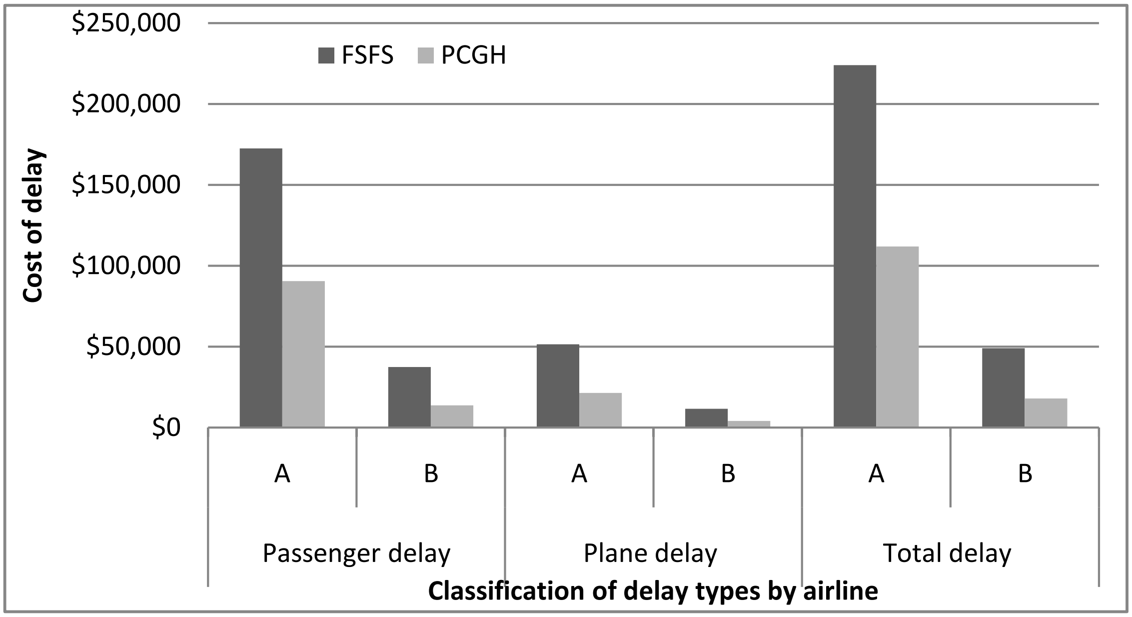

Figure 4-11: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a hub airport with one dominant airline with a runway capacity of 45 landings per hour, for the case when both airlines have the same distributions of times between connecting flights.

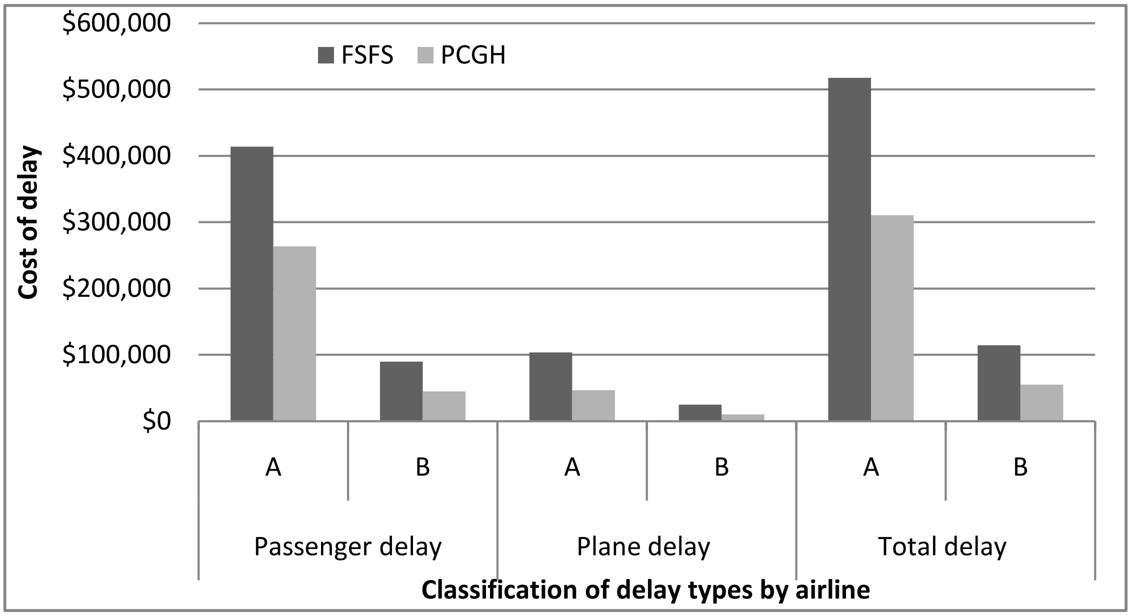

Figure 4-12: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a hub airport with one dominant airline with a runway capacity of 42 landings per hour, for the case when both airlines have the same distributions of times between connecting flights.

Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	Number of 10 minute time periods delayed							
				0	1	2	3	4	5	6	7
A	1	199	43.83	93	40	14	26	14	11	1	0
	2	202	2.50	187	15	0	0	0	0	0	0
	3	92	0.00	92	0	0	0	0	0	0	0
Total		493	46.33	372	55	14	26	14	11	1	0
B	1	54	12.33	18	14	14	4	1	2	1	0
	2	49	0.00	49	0	0	0	0	0	0	0
	3	18	0.00	18	0	0	0	0	0	0	0
Total		121	12.33	85	14	14	4	1	2	1	0

Table 4.13: Summary of delay for a PCGH allocation at a hub airport with one dominant airline with a runway capacity of 45 landings per hour, for the case when Airline B has longer times between connecting flights.

Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	Number of 10 minute time periods delayed												
				0	1	2	3	4	5	6	7	8	9	10	11	12
A	1	199	83.83	63	44	23	10	7	22	6	6	4	7	5	2	0
	2	202	8.83	149	53	0	0	0	0	0	0	0	0	0	0	0
	3	92	0.00	92	0	0	0	0	0	0	0	0	0	0	0	0
Total		493	92.67	304	97	23	10	7	22	6	6	4	7	5	2	0
B	1	54	24.33	9	10	9	8	7	4	3	4	0	0	0	0	0
	2	49	4.33	38	1	5	5	0	0	0	0	0	0	0	0	0
	3	18	0.00	18	0	0	0	0	0	0	0	0	0	0	0	0
Total		121	28.67	65	11	14	13	7	4	3	4	0	0	0	0	0

Table 4.14: Summary of delay for a PCGH allocation at a hub airport with one dominant airline with a runway capacity of 42 landings per hour, for the case when Airline B has longer times between connecting flights.

connections per amount of time delayed, while Airline A's passengers do. For Airline B, the distribution in Figure 4-2 is shifted 20 minutes to the right, which in turn causes the passenger delay cost functions for each aircraft type to decrease, making delays to Airline B less costly in the PCGH model.

PCGH Allocation

Tables 4.13 and 4.14 contain summaries of the delay statistics by airline for the PCGH allocations at runway capacities of 45 and 42 landings per hour, respectively. These tables can be contrasted with Tables 4.9 and 4.10 to see how the allocations are impacted by Airline B having longer times between connecting flights. The maximum delays of Airline B flights increase at both capacities. Notably, at capacity 42, Airline

B has Type 2 aircraft that are delayed for up to 30 minutes. In all previous PCGH allocations, the longest that a Type 2 aircraft is delayed is 10 minutes.

The difference in maximum delays of Type 2 aircraft illustrates the importance of marginal costs. A marginal cost is the cost of delaying a flight for an additional time period. The marginal cost for flight i at time period j is shown below in (4.1).

$$(C_{i,j} + D_{i,j}) - (C_{i,j-1} + D_{i,j-1}) \tag{4.1}$$

Since the passenger delay costs are a function of aircraft type and the number of passenger who are assumed to miss their connections at each time period, marginal costs are as well. For Airline A, the marginal cost of delaying a Type 1 aircraft for up to four time periods is less than that of delaying a Type 2 aircraft for one time period, but the marginal cost of delaying a Type 1 aircraft for five time periods is more. For Airline B, the marginal costs of delaying each aircraft type for one, two, or three time periods is the same, because no passengers miss connections until their flight is delayed for four time periods or more. This explains why Airline B has Type 2 aircraft that are delayed for up to 30 minutes at capacity 42; it is cheaper to delay an Airline B, Type 2 aircraft for these amounts of time than it is to delay an Airline A, Type 1 aircraft for 50 minutes or more. If Airline B had operated more flights during times of the day with the highest demand, more Type 2 aircraft would have been delayed.

Additionally, since Airline B's passengers begin missing connecting flights later than Airline A's, Airline B's total flight-hours of delay increase, as is expected. However, the magnitude of increase is not large, three and a half percent at capacity 45 and five percent at capacity 42. It is generalizable that airlines with longer times between connections are expected to be penalized more at higher congestion levels.

FSFS Allocation

Because the times at which connecting passengers are making connections had no impact on the FSFS allocations, Tables 4.11 and 4.12 from the previous subsection

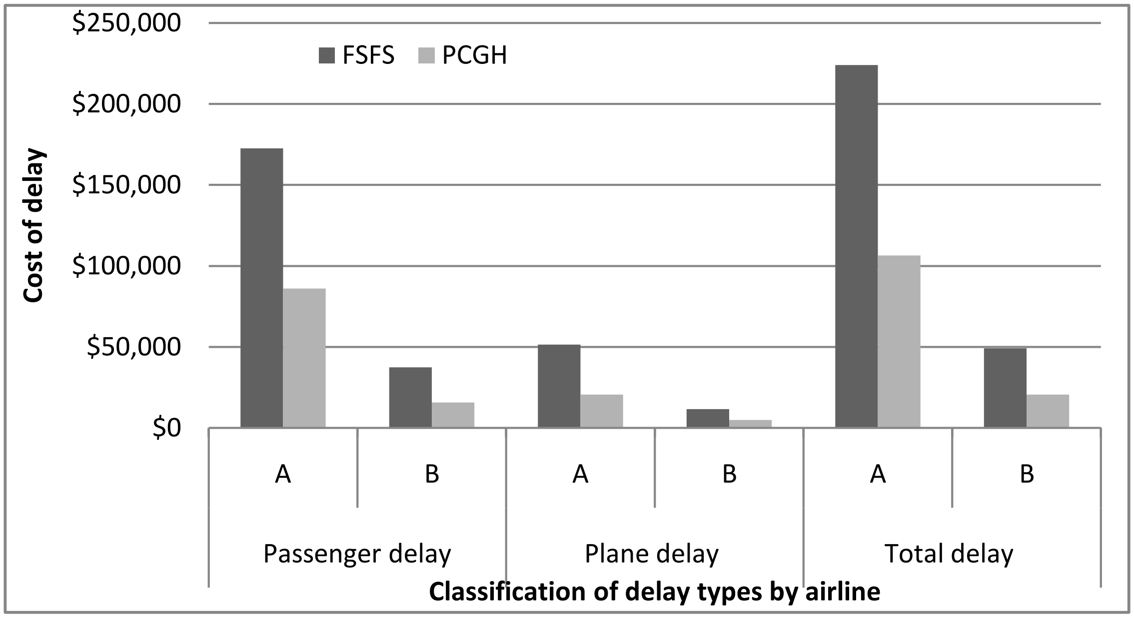

Figure 4-13: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a hub airport with one dominant airline with a runway capacity of 45 landings per hour, for the case when Airline B has longer times between connecting flights.

are still valid for this case.

PCGH vs. FSFS

Figures 4-13 and 4-14 contain comparisons of the delay costs associated with a PCGH vs. a FSFS allocation at a hub airport with one dominant airline with runway capacities of 45 and 42 landings per hour, respectively. These results are for the case when Airline B has longer times between connecting flights. As is expected since Airline B's total flight-hours of delay increase, Airline B's costs increase in all categories at both capacity levels when compared to the costs in Figures 4-11 and 4-12, while Airline A's decrease. The higher the congestion level, the more proportionally costly delay allocations are expected to be for an airline with longer times between connections.

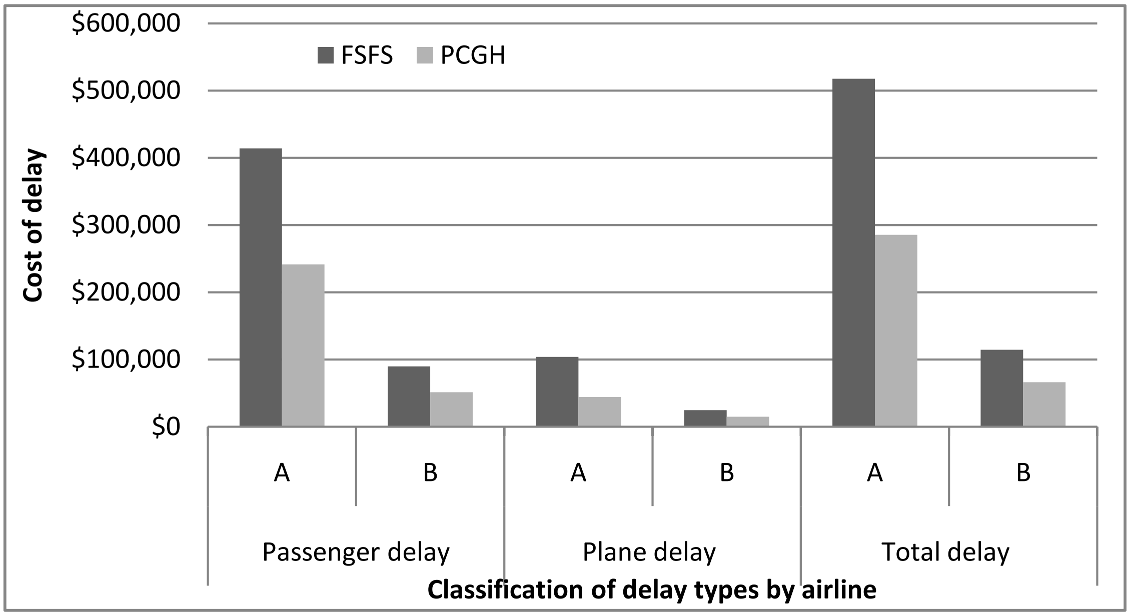

Figure 4-14: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a hub airport with one dominant airline with a runway capacity of 42 landings per hour, for the case when Airline B has longer times between connecting flights.

4.5 Hub Airport with Two Dominant Airlines

The third and final airport type presented is a hub with two dominant airlines. The types of airlines in this section are ones with majority and minority stakes in the airport, those with and without banks in their scheduling, and those with uniform fleets.

The assignment of airlines to flights is described in Section 4.5.1. Airlines A and B each are assumed to have approximately 40 percent of the flights in rotating two-hour banks, and Airline C is assumed to have the other 20 percent of flights, not in banks but randomly distributed throughout the day. Two assignments of airlines to flights are analyzed. In the first, each airline has the same fleet mix: 40 percent Type 1, 40 percent Type 2, and 20 percent Type 3 aircraft. In the second, Airline C has a fleet that is 100 percent Type 2, and Airlines A and B have fleet mixes of approximately 50 percent of Type 1 and 25 percent each of Types 2 and 3. For consistency, the assumption remains that 40 percent of the passengers on all flights

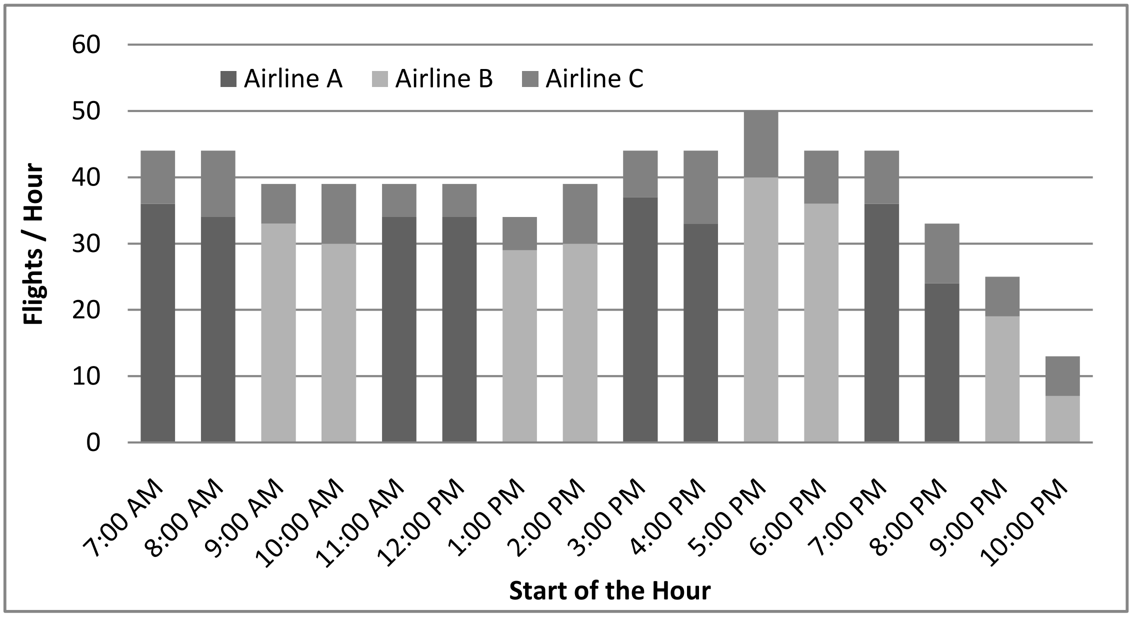

Figure 4-15: Hourly demand profile for Assignment 1 by airline for arrivals at a hub airport with two dominant airlines.

are making connections.

4.5.1 The Assignment of Airlines to Flights

Assignment 1: Equal Aircraft-Type Distributions per Airline

In the first assignment, each airline has the same fleet mix: 40 percent Type 1, 40 percent Type 2, and 20 percent Type 3 aircraft. To achieve this partitioning and the desired bank and non-bank schedules, first 20 percent of the flights throughout the day were randomly assigned to Airline C. Airlines A and B were then assigned to all of the remaining flights in alternating two-hour banks. Figure 4-15 shows the hourly demand profile by airline. The alternating two-hour banks of Airline A and B are easily seen, as are the non-bank Airline C assignments making up approximately 20 percent of each hour's demand.

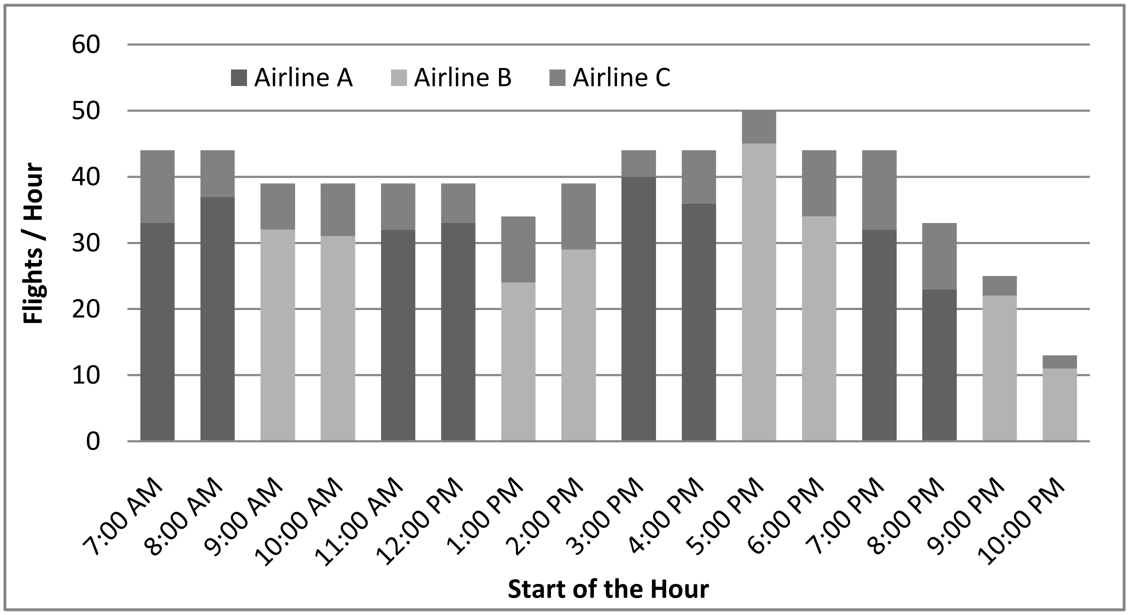

Figure 4-16: Hourly demand profile for Assignment 2 by airline for arrivals at a hub airport with two dominant airlines.

Assignment 2: Airline C with a Uniform Fleet of Type 2 Aircraft

In the second assignment, Airline C has a fleet that is 100 percent Type 2 aircraft. Since we are assuming the same deterministic arrival schedule with flight times and aircraft types already fixed, assigning Airline C to only Type 2 aircraft causes a change in the proportions of the flight types available to be assigned to Airlines A and B. In Assignment 2, we consider the case when Airlines A and B have the same fleet mix of approximately 50 percent Type 1 aircraft and 25 percent Type 2 and 3 aircraft. 20 percent of the Type 2 aircraft throughout the day were randomly assigned to Airline C. As in Assignment 1, Airlines A and B were then assigned to all of the remaining flights in alternating two hour banks. Figure 4-16 shows the hourly demand profile by airline. Though quite similar to Figure 4-15, the proportions of Airline C's flights in each hour differ in Figure 4-16, reflecting the change in Airline C's assignments.

4.5.2 Results for Assignment 1

The results for Assignment 1 for PCGH and FSFS are presented and contrasted in the following three subsections.

PCGH Allocation

Tables 4.15 and 4.16 contain summaries of the delay statistics by airline for the PCGH allocations at runway capacities of 45 and 42 landings per hour, respectively. These results are for Assignment 1, as described in Section 4.5.1, where all three airlines have the same fleet mixes. Total flight-hours of delay are approximately proportional to the number of flights operated per airline. However, Airline A receives its flight-hour delays from more Type 1 aircraft being delayed for less time than Airline B, which receives its flight-hour delays from delaying fewer flights for longer amounts of time. This is especially noticeable in Table 4.15 where 72 of Airline A's Type 1 flights are delayed compared to 43 of Airline B's. Also, in Table 4.16, it can be seen that both Airline B and C have flights delayed longer than eight time periods, while Airline A does not.

These differences in delay allocations result from the congestion levels during the times when the airlines operate. As can be seen in Figure 4-16, Airline A both operates the most flights and has the most flights delayed because their landings are during all of the highest demand two-hour banks except 5:00-7:00 p.m.. However, what is significant about the 5:00-7:00 p.m. bank is that it has the highest demand, and it is the block of time when the runway queue is the longest. The queue is at its maximum length because of the demand peak and the fact that demand consistently exceeds capacity for two hours prior. Because of this highest level of congestion, the airlines operating from 5:00-7:00 p.m., Airlines B and C, receive the longest delays of the day.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	112	23.00	40	34	19	11	7	1	0	0
	2	108	1.00	102	6	0	0	0	0	0	0
	3	48	0.00	48	0	0	0	0	0	0	0
Total		268	24.00	190	40	19	11	7	1	0	0
B	1	87	21.33	44	11	3	14	6	9	0	0
	2	96	0.67	92	4	0	0	0	0	0	0
	3	41	0.00	41	0	0	0	0	0	0	0
Total		224	22.00	177	15	3	14	6	9	0	0
C	1	54	11.83	26	8	7	7	3	2	1	0
	2	47	0.83	42	5	0	0	0	0	0	0
	3	21	0.00	21	0	0	0	0	0	0	0
Total		122	12.67	89	13	7	7	3	2	1	0

Table 4.15: Summary of delay for a PCGH allocation at a hub airport with two dominant airlines (Assignment 1) with a runway capacity of 45 landings per hour.

				Number of 10 minute time periods delayed												
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7	8	9	10	11	12
A	1	112	44.33	29	25	13	15	8	8	4	6	4	0	0	0	0
	2	108	5.00	78	30	0	0	0	0	0	0	0	0	0	0	0
	3	48	0.00	48	0	0	0	0	0	0	0	0	0	0	0	0
Total		268	49.33	155	55	13	15	8	8	4	6	4	0	0	0	0
B	1	87	42.67	28	16	10	5	1	11	1	0	4	4	5	2	0
	2	96	3.33	76	20	0	0	0	0	0	0	0	0	0	0	0
	3	41	0.00	41	0	0	0	0	0	0	0	0	0	0	0	0
Total		224	46.00	145	36	10	5	1	11	1	0	4	4	5	2	0
C	1	54	23.33	15	12	6	4	4	5	3	1	1	1	1	1	0
	2	47	2.67	31	16	0	0	0	0	0	0	0	0	0	0	0
	3	21	0.00	21	0	0	0	0	0	0	0	0	0	0	0	0
Total		122	26.00	67	28	6	4	4	5	3	1	1	1	1	1	0

Table 4.16: Summary of delay for a PCGH allocation at a hub airport with two dominant airlines (Assignment 1) with a runway capacity of 42 landings per hour.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	112	9.67	59	48	5	0	0	0	0	0
	2	108	10.50	50	53	5	0	0	0	0	0
	3	48	4.50	23	23	2	0	0	0	0	0
Total		268	24.67	132	124	12	0	0	0	0	0
B	1	87	7.50	55	19	13	0	0	0	0	0
	2	96	10.33	52	27	16	1	0	0	0	0
	3	41	4.83	20	14	6	1	0	0	0	0
Total		224	22.67	127	60	35	2	0	0	0	0
C	1	54	4.50	31	19	4	0	0	0	0	0
	2	47	3.67	27	18	2	0	0	0	0	0
	3	21	3.17	6	11	4	0	0	0	0	0
Total		122	11.33	64	48	10	0	0	0	0	0

Table 4.17: Summary of delay for a FSFS allocation at a hub airport with two dominant airlines (Assignment 1) with a runway capacity of 45 landings per hour.

FSFS Allocation

For Assignment 1, Tables 4.17 and 4.18 contain summaries of the delay statistics by airline for the FSFS allocations at runway capacities of 45 and 42 landings per hour, respectively. The delays, in flight-hours, are proportional to flights operated per airline. However, again Airline A receives more of its delay hours by delaying more flights for shorter amounts of time than Airlines B and C, as is the case in the PCGH allocation.

PCGH vs. FSFS

Figures 4-17 and 4-18 contain comparisons of the delay costs associated with a PCGH vs. a FSFS allocation at a hub airport with two dominant airlines with runway capacities of 45 and 42 landings per hour. These figures display the results for the first airline assignment, where each airline has a similar fleet mix. The airlines with majority stakes in the airport, Airlines A and B, have greater percentage improvements in cost savings from PCGH than Airline C at capacity 45; whereas percentage cost improvements were more varied across the airlines at capacity 42 (Airline A with the largest improvement, followed by Airline C, followed by Airline B).

It can be also be seen that Airline B has the highest passenger and total delay

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	112	19.83	45	35	16	12	4	0	0	0
	2	108	22.83	40	32	11	17	8	0	0	0
	3	48	9.17	18	15	5	10	0	0	0	0
Total		268	51.83	103	82	32	39	12	0	0	0
B	1	87	16.67	42	12	12	20	1	0	0	0
	2	96	18.83	41	22	10	21	2	0	0	0
	3	41	9.17	16	7	7	10	1	0	0	0
Total		224	44.67	99	41	29	51	4	0	0	0
C	1	54	8.83	22	18	8	5	1	0	0	0
	2	47	10.17	18	11	5	12	1	0	0	0
	3	21	5.83	4	7	4	4	2	0	0	0
Total		122	24.83	44	36	17	21	4	0	0	0

Table 4.18: Summary of delay for a FSFS allocation at a hub airport with two dominant airlines (Assignment 1) with a runway capacity of 42 landings per hour.

costs at capacity 45. Yet Airline A has the most flight-delay hours for Type 1 and 2 aircraft. This disparity illustrates that it is not the amount of total delay but *which type of aircraft is delayed and for how long* that impacts passenger costs. For Airline B, operating more flights during the most congested times of the day resulted in more missed connections and higher passenger delay costs, which drove the total delay costs higher.

4.5.3 Results for Assignment 2

Results for Assignment 2 for PCGH and FSFS are presented in the following three subsections. Recall that in Assignment 2, Airline C has a fleet that is 100 percent Type 2, and Airlines A and B have fleet mixes of approximately 50 percent of Type 1 and 25 percent each of Types 2 and 3 (see Assignment 2 in Section 4.5.1).

PCGH Allocation

Tables 4.19 and 4.20 contain summaries of the delay statistics by airline for the PCGH allocations at runway capacities of 45 and 42 landings per hour, respectively. In both tables, the differences in the distributions of delay among the airlines are stark when compared with those in Tables 4.15 and 4.16. Airline C receives less than five percent

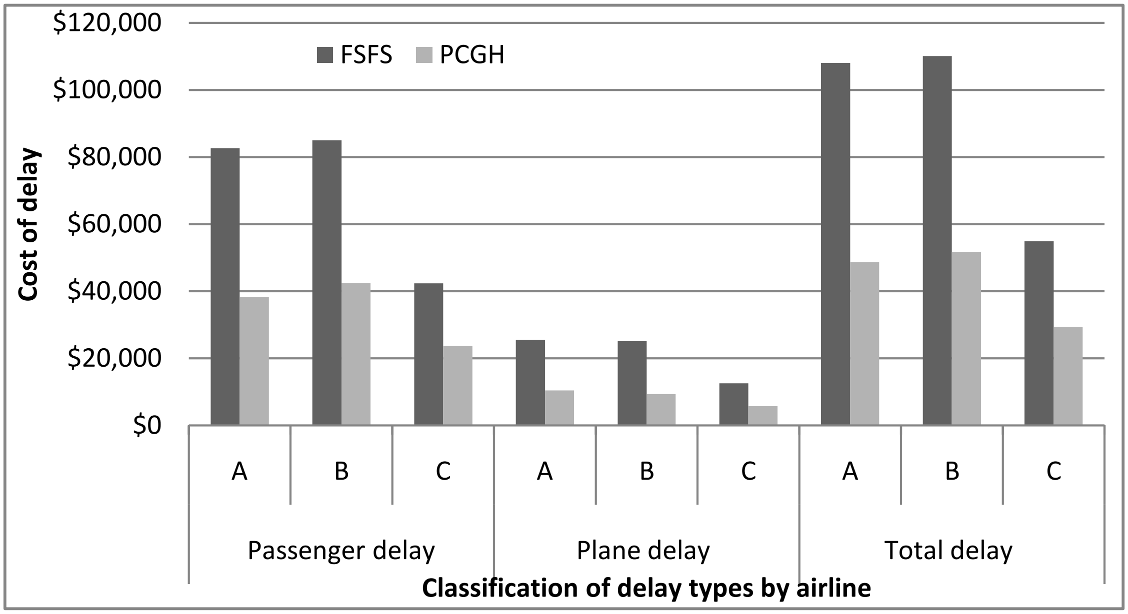

Figure 4-17: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a hub airport with two dominant airlines (Assignment 1) with a runway capacity of 45 landings per hour.

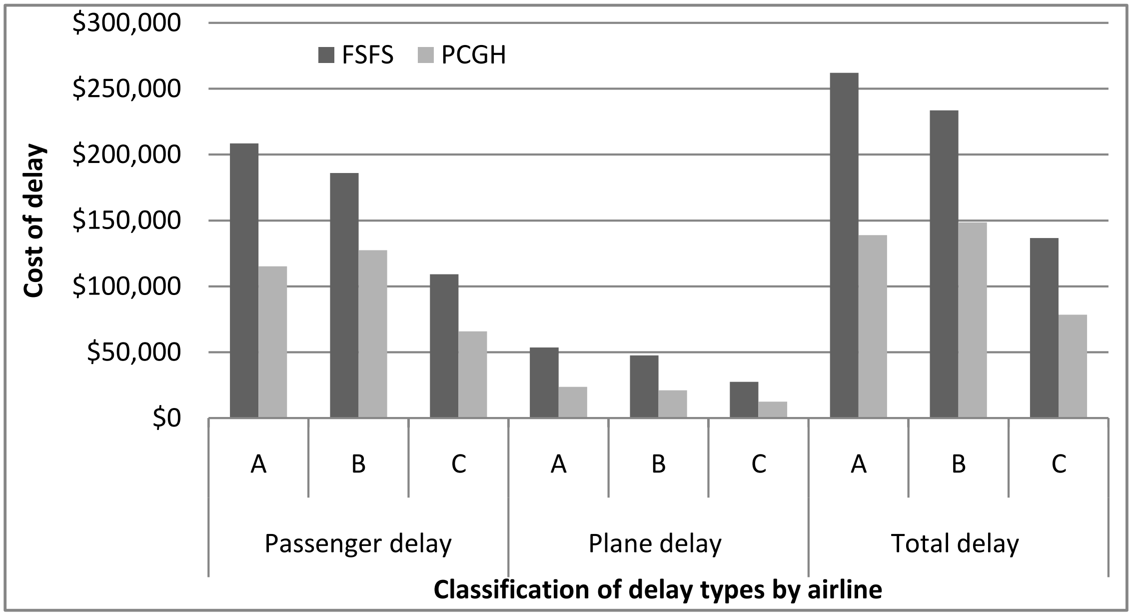

Figure 4-18: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a hub airport with two dominant airlines (Assignment 1) with a runway capacity of 42 landings per hour.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	137	28.67	49	39	24	16	8	1	0	0
	2	68	1.00	62	6	0	0	0	0	0	0
	3	61	0.00	61	0	0	0	0	0	0	0
Total		266	29.67	172	45	24	16	8	1	0	0
B	1	116	27.50	61	14	5	16	8	11	1	0
	2	63	0.83	58	5	0	0	0	0	0	0
	3	49	0.00	49	0	0	0	0	0	0	0
Total		228	28.33	168	19	5	16	8	11	1	0
C	1	0	0.00	0	0	0	0	0	0	0	0
	2	120	0.67	116	4	0	0	0	0	0	0
	3	0	0.00	0	0	0	0	0	0	0	0
Total		120	0.67	116	4	0	0	0	0	0	0

Table 4.19: Summary of delay for a PCGH allocation at a hub airport with two dominant airlines and a third with a fleet of only Type 2 aircraft (Assignment 2) with a runway capacity of 45 landings per hour.

of the flight-hours of delay at both capacity levels, despite operating 20 percent of the flights during the day. This is not surprising since Type 2 aircraft received only small delays from the PCGH allocations, and Airline C operates only Type 2 aircraft.

FSFS Allocation

For Assignment 2, Tables 4.21 and 4.22 contain summaries of the delay statistics by airline for the FSFS allocations at runway capacities of 45 and 42 landings per hour, respectively. In this allocation, the flight-hours of delay are allocated equitably among the airlines, as is expected since the FSFS model does not discriminate based on aircraft type.

PCGH vs. FSFS

Figures 4-19 and 4-20 contain comparisons of the delay costs associated with a PCGH vs. a FSFS allocation at a hub airport with two dominant airlines with runway capacities of 45 and 42 landings per hour, respectively. These figures display the results for the second airline assignment, where Airline C has a fleet that is 100 percent Type 2, and Airlines A and B have fleet mixes of approximately 50 percent

				Number of 10 minute time periods delayed												
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7	8	9	10	11	12
A	1	137	55.67	32	32	16	19	12	8	6	7	5	0	0	0	0
	2	68	2.83	51	17	0	0	0	0	0	0	0	0	0	0	0
	3	61	0.00	61	0	0	0	0	0	0	0	0	0	0	0	0
Total		266	58.50	144	49	16	19	12	8	6	7	5	0	0	0	0
B	1	116	54.67	40	21	13	5	1	16	2	0	4	5	6	3	0
	2	63	2.67	47	16	0	0	0	0	0	0	0	0	0	0	0
	3	49	0.00	49	0	0	0	0	0	0	0	0	0	0	0	0
Total		228	57.33	136	37	13	5	1	16	2	0	4	5	6	3	0
C	1	0	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
	2	120	5.50	87	33	0	0	0	0	0	0	0	0	0	0	0
	3	0	0.00	0	0	0	0	0	0	0	0	0	0	0	0	0
Total		120	5.50	87	33	0	0	0	0	0	0	0	0	0	0	0

Table 4.20: Summary of delay for a PCGH allocation at a hub airport with two dominant airlines and a third with a fleet of only Type 2 aircraft (Assignment 2) with a runway capacity of 42 landings per hour.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	137	12.00	71	60	6	0	0	0	0	0
	2	68	6.50	32	33	3	0	0	0	0	0
	3	61	5.83	28	31	2	0	0	0	0	0
Total		266	24.33	131	124	11	0	0	0	0	0
B	1	116	9.67	74	26	16	0	0	0	0	0
	2	63	6.00	35	20	8	0	0	0	0	0
	3	49	6.67	21	17	10	1	0	0	0	0
Total		228	22.33	130	63	34	1	0	0	0	0
C	1	0	0.00	0	0	0	0	0	0	0	0
	2	120	12.00	62	45	12	1	0	0	0	0
	3	0	0.00	0	0	0	0	0	0	0	0
Total		120	12.00	62	45	12	1	0	0	0	0

Table 4.21: Summary of delay for a FSFS allocation at a hub airport with two dominant airlines and a third with a fleet of only Type 2 aircraft (Assignment 2) with a runway capacity of 45 landings per hour.

				Number of 10 minute time periods delayed							
Airline	Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7
A	1	137	24.83	51	48	18	15	5	0	0	0
	2	68	13.33	25	22	7	12	2	0	0	0
	3	61	12.00	21	20	8	12	0	0	0	0
Total		266	50.17	97	90	33	39	7	0	0	0
B	1	116	20.50	58	17	18	22	1	0	0	0
	2	63	12.00	27	13	10	13	0	0	0	0
	3	49	12.17	17	9	8	12	3	0	0	0
Total		228	44.67	102	39	36	47	4	0	0	0
C	1	0	0.00	0	0	0	0	0	0	0	0
	2	120	26.50	47	30	9	25	9	0	0	0
	3	0	0.00	0	0	0	0	0	0	0	0
Total		120	26.50	47	30	9	25	9	0	0	0

Table 4.22: Summary of delay for a FSFS allocation at a hub airport with two dominant airlines and a third with a fleet of only Type 2 aircraft (Assignment 2) with a runway capacity of 42 landings per hour.

of Type 1 and 25 percent of Type 2 and 3 aircraft. It is easily seen that Airline C benefits much more from PCGH than Airline A and B in cost savings in all categories and more so at capacity 45 than 42. Since PCGH primarily assigns delays to Type 1 aircraft and Airline C does not operate any Type 1 aircraft, these benefits are not surprising. It also makes sense that the percentage savings are higher for Airline C at capacity 45 than at capacity 42 because, in PCGH, the more the capacity is reduced the more likely delays are to be allocated to Type 2 aircraft. The generalization is that delay allocations to airlines in PCGH are strongly linked to fleet mixes. The fewer Type 1 aircraft operated, the less delay; and at higher congestion levels, the fewer Type 2 aircraft, the less delay.

4.6 Results Using the Maximum Delay Constraint

One potential criticism of the PCGH allocations is that some of the delays that are assigned to aircraft with fewer passengers are very long. In this section, we address the impact of reducing the length of the longest delays by using the maximum delay constraint. Modified PCGH allocations are presented as are the resulting relative increases in delay costs. A noteworthy result is that, even with the maximum delay

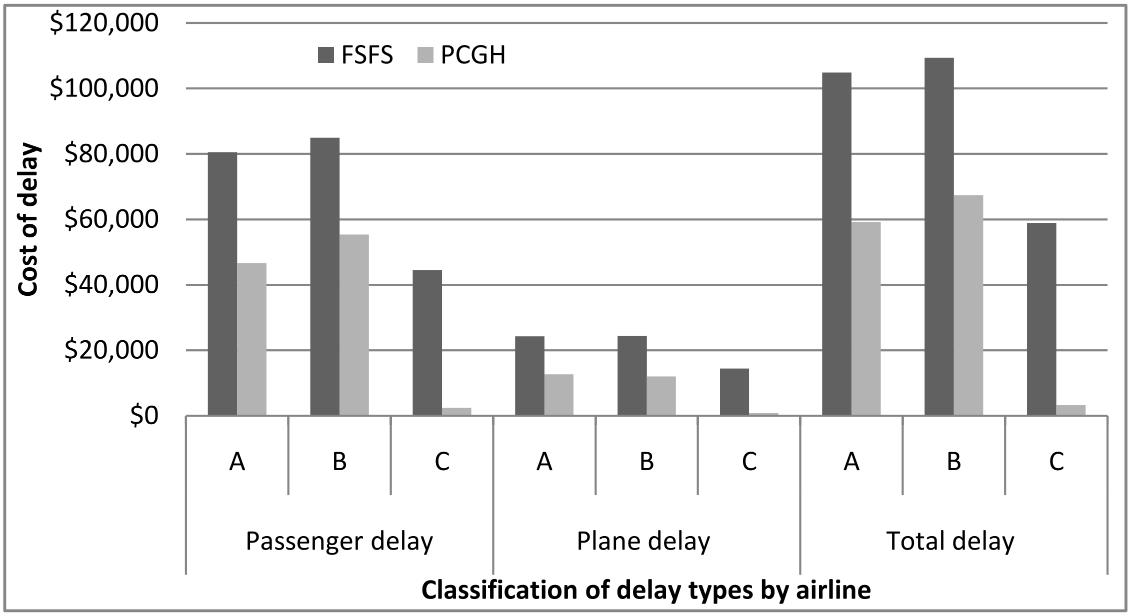

Figure 4-19: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a hub airport with two dominant airlines and a third with a fleet of only Type 2 aircraft (Assignment 2) with a runway capacity of 45 landings per hour.

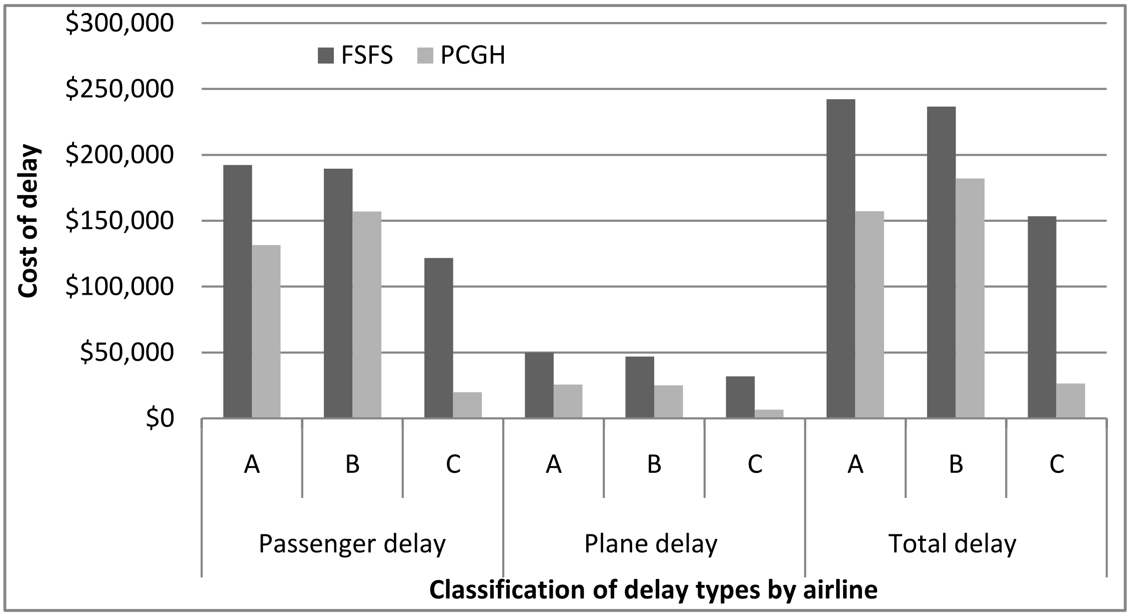

Figure 4-20: A comparison of the delay costs associated with a FSFS allocation vs. a PCGH allocation at a hub airport with two dominant airlines and a third with a fleet of only Type 2 aircraft (Assignment 2) with a runway capacity of 42 landings per hour.

constraint, PCGH still reduced delay costs when compared to FSFS.

As detailed in Section 3.3.1, to address the issue of exceedingly long delays for flights with fewer passengers, the "flights eventually land" constraint in the PCGH formulation, (3.6), can be modified to limit the length of delays for individual flights. The modified constraint (4.2) is shown below. Each flight i is assigned a parameter, M_i, that specifies the maximum number of time periods by which the flight may be delayed.

$$\sum_{j=P_i}^{P_i+M_i} x_{ij} = 1, \quad \forall i \in \{1, ..., N\} \tag{4.2}$$

For a runway capacity of 42 landings per time period, the PCGH model was run with the maximum allowable delay to flights set at three different and decreasing levels. The delay allocation results are summarized in Table 4.23. The first and second levels were $M_i = 8$ and $M_i = 4$, respectively, for all flights i. Setting M_i lower than 4 for all flights made the problem infeasible. This was expected since delays of 4 time periods do occur in the FSFS allocation at capacity 42. However, the longest delays in the FSFS allocation at the non-peak times of the day were less than 4 time periods, and the third level of maximum allowable delay took advantage of this. The third delay level was $M_i = 2$ for flights scheduled to land from 7:00 a.m to 2:49 p.m., $M_i = 3$ from 2:50 p.m. to 4:49 p.m., and $M_i = 4$ from 4:50 p.m. to 10:59 p.m.. In other words, at this third level, M_i was set to the maximum instance of FSFS delay for each block of time. To make comparisons easier, the results for FSFS and PCGH allocations without the maximum delay constraint from Section 4.2 are included at the top of the table. As M_i decreases, the delays to Type 1 flights decrease, while delays to Type 2 and 3 flights increased to compensate, causing the allocations to become closer to FSFS.

At the second and third levels of delay (with $M_i = 4$ during the peak hours), an apparent anomaly is that more Type 3 aircraft are delayed for two time periods than for one. However, a comparison of marginal costs explains these assignments. The Type 3 flights delayed for two time periods are those scheduled to land at the very

			Number of ten minute time periods delayed												
Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7	8	9	10	11	12
FSFS															
1	253	45.33	109	65	36	37	6	0	0	0	0	0	0	0	0
2	251	51.83	99	65	26	50	11	0	0	0	0	0	0	0	0
3	110	24.17	38	29	16	24	3	0	0	0	0	0	0	0	0
Total	614	121.33	246	159	78	111	20	0	0	0	0	0	0	0	0
PCGH with $M_i = \infty$ (technically $M_i = P + 1 - P_i$)															
1	253	110.33	72	53	29	24	13	24	8	7	9	5	6	3	0
2	251	11	185	66	0	0	0	0	0	0	0	0	0	0	0
3	110	0	110	0	0	0	0	0	0	0	0	0	0	0	0
Total	614	121.33	367	119	29	24	13	24	8	7	9	5	6	3	0
PCGH with $M_i = 8$															
1	253	106	72	53	29	24	13	24	8	7	23	0	0	0	0
2	251	15.33	184	45	19	3	0	0	0	0	0	0	0	0	0
3	110	0	110	0	0	0	0	0	0	0	0	0	0	0	0
Total	614	121.33	366	98	48	27	13	24	8	7	23	0	0	0	0
PCGH with $M_i = 4$															
1	253	80.17	72	53	29	26	73	0	0	0	0	0	0	0	0
2	251	36	170	16	12	36	17	0	0	0	0	0	0	0	0
3	110	5.17	94	1	15	0	0	0	0	0	0	0	0	0	0
Total	614	121.33	336	70	56	62	90	0	0	0	0	0	0	0	0
PCGH with M_i *= 2 from 7:00am - 2:49pm, 3 from 2:50pm - 4:49pm, and 4 from 4:50pm-10:59pm*															
1	253	77.33	72	53	32	37	59	0	0	0	0	0	0	0	0
2	251	38.83	157	25	16	36	17	0	0	0	0	0	0	0	0
3	110	5.17	94	1	15	0	0	0	0	0	0	0	0	0	0
Total	614	121.33	323	79	63	73	76	0	0	0	0	0	0	0	0

Table 4.23: Summary of PCGH delay allocations using the maximum delay constraint for a runway capacity of 42 landings per hour. For reference, allocations for FSFS and PCGH without limitations on maximum delay are included.

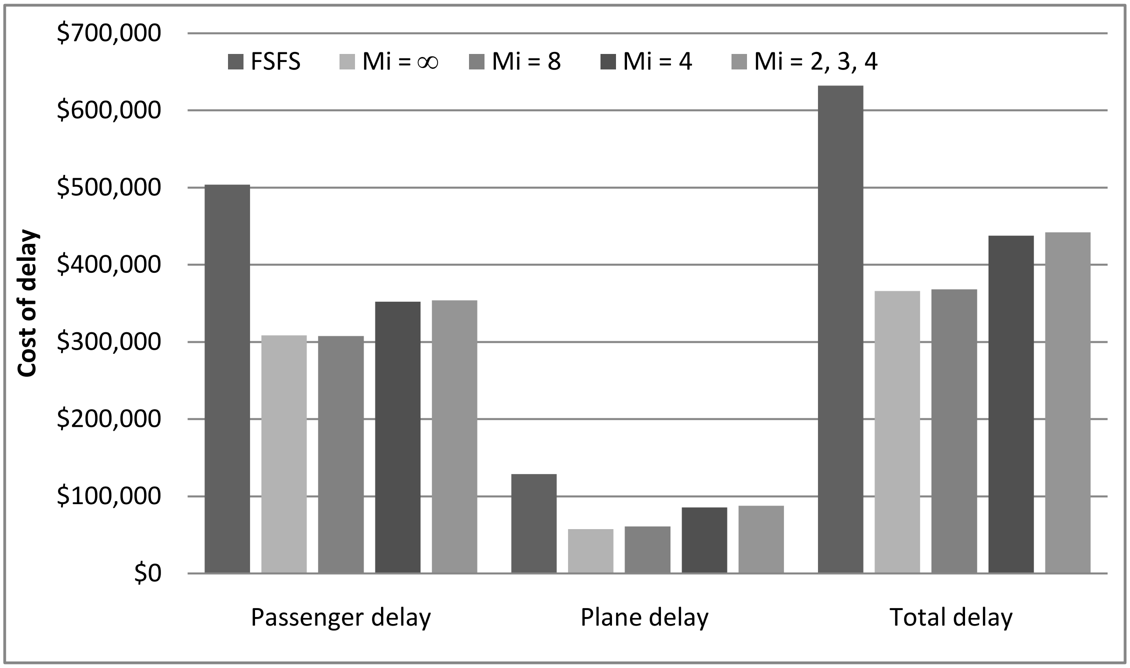

Figure 4-21: A comparison of delay costs for FSFS and PCGH with various levels of maximum delay for a runway capacity of 42 landings per hour.

busiest times of the day, during which all of the Type 1 flights are delayed to their limit of four time periods and the rest of the necessary delay must be assigned to Type 2 and 3 aircraft. The marginal cost of delaying a Type 3 aircraft for two time periods is less than that of delaying a Type 2 aircraft for four time periods (recall from Figure 4-2 that a delay of 40 minutes is assumed to cause a large spike in the number of passengers missing their connections), while the marginal cost of delaying a Type 2 aircraft for four time periods is less than that of delaying a Type 3 aircraft for three time periods. Delay assignments are thus made accordingly in a balancing act of marginal costs. Once as many Type 2 aircraft as possible have been delayed for three time periods, Type 3 aircraft are delayed for up to two periods. The remaining necessary delay is then absorbed by increasing the delay of some of the Type 2 aircraft from three to four time periods.

While the delay costs increase each time the maximum delay limit is decreased (due to more delays to Type 2 and 3 aircraft), the total delay costs are still much smaller than those for the FSFS allocation. Figure 4-21 shows a comparison of delay costs for the allocations in Table 4.23. For all settings of M_i in our experiments, the increase in PCGH total costs is 20 percent at most, and the percentage savings over FSFS remain at least 30 percent. Though the savings of PCGH over FSFS could have been reduced to zero if, for every flight, M_i was set to the number of time periods that flight i was delayed in the FSFS allocation, these results indicate that much of the savings of PCGH can be achieved without assigning such lengthy delays to smaller flights.

4.7 Dependence of the Results on the Convexity of the Cost Function

This final section of Chapter 4 addresses the dependence of the results of the computational tests on the convexity of the cost function. We present an alternative cost function that is non-convex and that approximates passenger delay costs for missing connections in a more logically consistent manner (always under the assumption of deterministic knowledge of what future flights will have seats available to passengers' final destinations). For this function, it was seen that, in the absence of convexity of the delay cost functions, Type 1 aircraft incurred even longer delays than with the original PCGH cost functions. Moreover, these delays often display an interesting pattern.

An alternative, non-convex passenger delay cost function for flight i, D_{ij}, is shown in Figure 4-22 for Type 1, Type 2, and Type 3 aircraft. This function will be referred to, in a colloquial way, as the *convex-concave cost function*, while the cost function that was used in our earlier analysis as the *PCGH cost function*. Both cost functions can be thought of as "passenger-centric." The convex-concave cost function takes a convex shape during the times when passengers are missing their flights and a concave shape between these times. As in the PCGH cost function, the cost of delay for a flight in the convex-concave cost function is obtained by summing the costs of the individual delay functions of all of the passengers onboard. For non-connecting

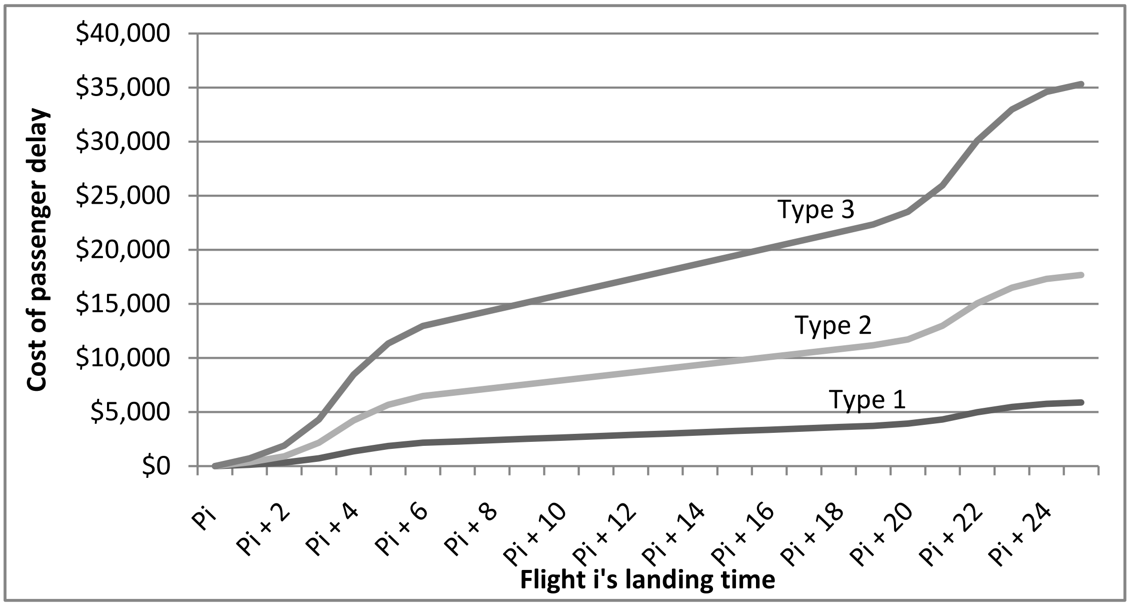

Figure 4-22: Convex-concave passenger delay costs for some flight i, if i is Type 1, Type 2, or Type 3.

passengers, these individual cost functions are the same as in the PCGH cost function in our computational tests, with a constant slope of $30 per hour. For connecting passengers, the structures of the individual cost functions are different from the PCGH case, although the distributions of the times when passengers are assumed to miss their scheduled connections are the same for each aircraft type as in the PCGH cost function (see Figure 4-2).

The structure of the individual cost functions for connecting passengers in the convex-concave cost function is shown in Figure 4-23 and should be compared to the PCGH individual cost function structure in Figure 3-2. The function is flat except for the large step increases at flight-missing time periods. It is assumed that the amount of delay that would cause a passenger to miss his or her original connection is known. Similarly, the amount of delay that would cause the passenger to miss subsequent flights with available seats to the passenger's final destination is also assumed known. For example, in the functions shown in Figure 4-22, flights with open seats to connection destinations are assumed to depart every three hours from airport Z.

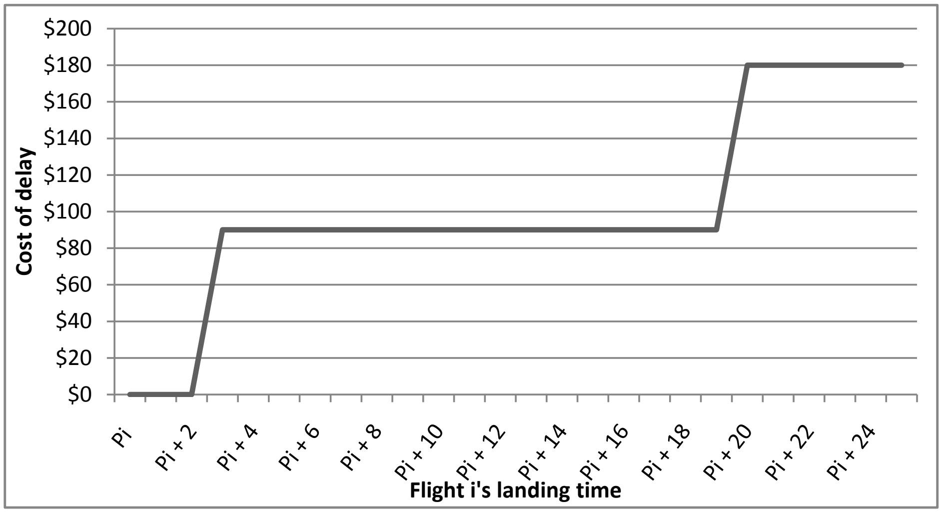

Figure 4-23: Individual delay costs for a connecting passenger in the convex-concave cost function who is assumed to miss his/her connecting flight at three time periods of delay, with the next connecting flight three hours later.

The deterministic knowledge of future flights makes the costs in the convex-concave function more accurate in the sense that the steps reflect the cost of how much later a passenger will arrive at his/her destination due to missing subsequent each flight. The cost of each step (its height) is the amount of time until the next flight-missing step multiplied by the standard rate of $30 per hour. In Figure 4-23, the first step occurs at a delay of three time periods, which is when the passenger is assumed to miss his or her original connection. Since the next flight to the passenger's final destination is assumed to be three hours later than the original connection, the height of the step is $90. The reason that the function is flat in-between steps is because arriving earlier or later between steps does not impact the time of arrival to the passenger's final destination.

Table 4.24 shows the results of using the convex-concave cost function for a runway capacity of 42 landings per hour. In the allocation, delays to Type 1 flights are even longer and more disproportionate than in the allocation using the PCGH cost function in Table 4.2. The most noticeable trait of the convex-concave allocation is that Type 1 flights are either delayed for thirty minutes or less or for over two hours, with no gradation in between. This follows directly from the structure of the convex-concave

			Number of 10 minute time periods delayed																					
Flight Type	# Flights	Tot. Flight-Hr Delay	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
1	253	119.33	76	88	55	7	0	0	0	0	0	0	0	0	0	0	1	2	1	2	2	13	6	0
2	251	2.00	239	12	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	110	0.00	110	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Total	614	121.33	425	100	55	7	0	0	0	0	0	0	0	0	0	0	1	2	1	2	2	13	6	0

Table 4.24: Summary of delay for a passenger-centric allocation using the convex-concave cost function for a runway capacity of 42 landings per hour.

cost function: convex before four time periods of delay and then concave afterwards for two and half hours. Because convexity and concavity are directly linked to whether marginal costs are increasing or decreasing, this means that marginal costs increase from zero to four time periods of delay and then decrease or stay the same afterwards. For this reason, once a flight is delayed for four time periods, it is cheaper to continue to delay the flight for longer amounts of time than it is to delay another flight for two to four time periods. Only when the connecting passengers on a delayed flight begin reaching the second step in their individual cost functions do the marginal costs begin increasing again, giving incentive to stop adding length to the delay of the flight. It is hard to imagine widespread support for a ground holding program that so heavily penalizes a small percentage of flights.

While it is not necessarily intuitive that the slope in the PCGH individual cost function stays the same for all time periods after the flight-missing breakpoint, if this slope became less steep at some point in time, the function would no longer be convex, causing the same problem of gaps in delay lengths as occurs in the convex-concave cost function allocation. In addition to it being unattractive to delay some flights for much longer than others, delays in the NATS are often too interconnected to accurately predict the time of the next connecting flight for a passenger. Thus, the PCGH cost function is attractive in the sense that it models the steep increase in delay costs of missing connections without requiring exact knowledge of when the next seats will be available on flights to passengers' final destinations. Further, if such information were available, it could be reflected in the steepness of the slope for individual PCGH cost functions; steeper if the connection was later and less steep if it was sooner. On the negative side, a disadvantage of using the PCGH cost function is that the passenger delay costs are not as intuitively appealing as those in the convex-concave cost function. The costs are approximations at best, but the increasing marginal costs prevent an aircraft from being delayed just because most of its connecting passengers have already missed their connections.

Chapter 5

Conclusions and Further Research

5.1 Conclusions

The results of our research indicate potential for improving passenger delay costs and overall delay costs by using a passenger-centric allocation of delays in GDPs. We proposed a Passenger-Centric Ground Holding (PCGH) model, based on the Terrab-Odoni deterministic integer program from [12], which took into consideration, with appropriately scaled costs, both the number of passengers on flights and when/if they were connecting. In the particular numerical examples presented in Section 4.2, PCGH was shown to save an average of 45 percent over FSFS in passenger, aircraft, and total costs. The exact magnitude of savings will, of course, vary depending on the case at hand. The analysis presented here has many limitations and is based on a number of simplifying assumptions. However, our results suggest that, under certain circumstances, the savings obtained through a PCGH-type of model may be significant. It is therefore worthwhile to examine passenger-centric GDP strategies in more depth, despite the obvious implementation difficulties that such strategies may have to address in practice.

As a whole and by design, PCGH delayed fewer flights than did FSFS. While most Type 2 and 3 aircraft saw shorter delays in PCGH than they would have in FSFS, many Type 1 aircraft saw longer delays. The extent to which these latter delays were longer depended on the prevalent level of congestion. The less congestion, the closer the PCGH allocations were to the FSFS ones, in the sense that the longest delay assigned in PCGH was closer to the longest delay assigned in FSFS. However, at lower congestion levels, the delay was served almost exclusively by Type 1 aircraft, resulting in higher percentages of savings, at the expense of greater imbalance in the types of aircraft being delayed.

What can be generalized from the airport and airline-specific results in Sections 4.3 - 4.5 is that PCGH allocations treat equitably, in terms of proportions of delay, airlines with similar percentages of connecting passengers, times between connecting flights, and fleets. However, this equity is highly dependent on how many aircraft of each type are operated by each airline during the peaks and valleys in the airport demand, especially for airlines using banks. While the lengths of delays in the FSFS allocations also depend on how many flights each airline operates during the peaks and valleys in demand, they do not depend on the aircraft type. The difference in aircraft-type dependence between PCGH and FSFS led to variations in the percentage improvements in costs achieved by different airlines, especially in the non-hub case.

As expected, when airlines had different times between connections or different fleets, the results were not as equitable. In Section 4.4, the delays were longer and more costly for the airline with the longer scheduled times between arrival and departure banks. More generally, it was observed that the higher the congestion level, the more the airline with the longer times between connections was penalized. In Section 4.5, the delays were shorter and less costly for the airline with a uniform fleet of Type 2 aircraft. The generalized observation is that delay allocations to airlines in PCGH are strongly linked to fleet mixes. The fewer Type 1 aircraft operated, the less delay; and at higher congestion levels, the fewer Type 2 aircraft, the less delay. A somewhat surprising finding was that fleet variations had a much bigger impact than the lengths of times between connecting flights on the equitable treatment of the airlines in PCGH. In hindsight, this follows from the fact that the differences in marginal costs among aircraft types were so large.

In addition, the effects of the maximum-delay-limiting constraint and the convexity of the cost function were considered in Sections 4.6 and 4.7, respectively. The maximum delay constraint impacted the extent of the differences between the PCGH and the FSFS allocations. As M_i, the maximum delay allowed, decreased, the PCGH allocations became closer to the FSFS allocations with more Type 2 and 3 flights delayed and shorter delays for Type 1 flights. However, even at the lowest feasible settings of maximal delay by time period, cost savings over FSFS were still high in our computational tests, indicating that most of the savings of PCGH could still be achieved without assigning excessive delays to smaller flights. Last, it was seen that, in the absence of convexity of the cost function, Type 1 flights incurred even longer and unevenly distributed delays.

Although these results are quite promising, it is worth emphasizing again that the schedule, aircraft types, and airline assignments used in this thesis were based on a hypothetical example and, therefore, their implications require further investigation. For the sake of comparison and consistency, we used the same airport arrival profile and demand rates for the three types of airports examined. What is important is asking the right questions, and that is what we tried to do in a systematic manner. With a real schedule for an actual hub or non-hub airport, it would not be difficult to use the same type of analysis to produce results with more specific implications. Likewise, with a more theoretical approach, more generalizable results could be obtained. We believe that future research in making GDPs more passenger-centric would be worthwhile.

5.2 Further Research

The PCGH model could be improved by more realistically incorporating the multiple types of flows through the air transportation network, namely the aircraft, crews, and passengers. As an example, consider some arriving Flight A where all passengers are connecting to some other departing Flight C that is scheduled to leave 30 minutes after Flight A is scheduled to arrive. Neglecting, without loss of generality, travel time within the airport, the flight-missing breakpoint for all passengers on Flight A occurs 29 minutes after Flight A's scheduled arrival time, so to minimize cost, the

PCGH model assigns Flight A a landing slot before the breakpoint. However, the missing detail may be that departing Flight C is operated by the same aircraft as arriving Flight B, and arriving Flight B could have relatively few connecting passengers onboard and thus be delayed one or two hours. This would mean that the breakpoint for Flight A was really much later than in the original parameter assignment. The exact same issue could arise if, instead of the aircraft, the crew of the Flight C were arriving on Flight B.

The air transportation network is enormously complex and interconnected. A natural extension of the PCGH model would be to try to reflect, in each individual connecting passenger cost function, the dependence of the flight-missing breakpoint on the arrival times of a set of flights, including more than just the one the passenger is on. Additionally, many other models in the literature, including the network-wide formulations in [4] and [8], use flight-specific costs. The PCGH cost function could easily be adapted for incorporation into these models.

A more far-reaching issue concerns the effectiveness of the current CDM-based method of decentralized airline decision-making in allocating flights to arrival slots during GDPs. In the current system, once airlines receive their assigned slots from the FAA in a FSFS manner, known as Ration by Schedule (RBS), they are free to make swaps and change the ordering of the arrivals of their flights. In this process, they can address, to the extent that they see fit, the issue of the number of passengers on each flight and, if applicable, their corresponding connection times. Under this structure, passengers are only indirectly represented. With the decentralized allocation approach, the number of landing slots that each airline has available for swapping is limited to those which they received from the RBS allocation. In this sense, a decentralized allocation that is optimized for passengers by each airline is a lower bound on the passenger-optimal swaps that a centralized approach could make. In other words, if the FAA could violate RBS scheduling for the sake of minimizing the total passenger delay cost in the system, this centralized allocation would be no worse than the resulting schedule from the best-possible scenario of decentralized swaps that the airlines could make on behalf of passengers.

An interesting research question to investigate is the extent to which the optimal decentralized versus optimal centralized ground holding allocations are different from a passenger's perspective. As mentioned in Section 2, Hanowsky did this type of analysis with his passenger-driven cost functions that did not consider connections. For the airport he was analyzing, he found significant improvements from using a centralized instead of decentralized approach, [7]. This type of experiment could be replicated using the PCGH cost function.

Bibliography

[1] Total delay impact: Interim report. Technical report, Brattle Group, 2009.

[2] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications.* Prentice Hall, 1993.

[3] C. Barnhart, D. Bertsimas, C. Caramanis, A.R. Odoni, and G. Perakis. The national air transportation as a reconfigurable engineered system (nsf-efri proposal), 2007.

[4] Dimitris Bertsimas and Sarah Stock Patterson. The air traffic flow management problem with enroute capacities. *Operations Research*, 46(3):406–422, 1998.

[5] Dimitris Bertsimas and John N. Tsitsiklis. *Introduction to Linear Optimization.* Athena Scientific, 1997.

[6] Stephane Bratu and Cynthia Barnhart. Flight operations recovery: New approaches considering passenger recovery. *J. of Scheduling*, 9(3):279–298, 2006.

[7] Michael J. Hanowsky. *A Model to Design a Stochastic and Dynamic Ground Delay Program Subject to Non-Linear Cost Functions.* PhD thesis, Massachusetts Institute of Technology, 2008.

[8] A. R. Odoni Lulli, G. The european air traffic flow management problem. *Transportation Science*, 2007.

[9] A.R. Odoni. The flow management problem in air traffic control. *Flow Control of Congested Networks*, pages 269–288, 1987.

[10] Lance Sherry and George Donohue. U.s. airline passenger trip delay report, 2007. `http://catsr.ite.gmu.edu`.

[11] Mostafa Terrab. *Ground Holding Strategies for Air Traffic Control.* PhD thesis, Massachusetts Institute of Technology, 1990.

[12] Mostafa Terrab and Amedeo R. Odoni. Strategic flow management for air traffic control. *Operations Research*, 41(1):138–152, 1993.

[13] P.B. Vranas, D.J. Bertsimas, and A.R. Odoni. The multi-airport ground-holding problem in air traffic control. *Operations Research*, 42(2), 1994.