Analysis of Efficiency and Fairness in Stochastic Ground Holding Models

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
JUN 28 2016
LIBRARIES
ARCHIVES

by

Karthik Kavassery Gopalakrishnan

B.Tech., Indian Institute of Technology Madras (2014)

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016

© Massachusetts Institute of Technology 2016. All rights reserved.

Signature redacted

Author .....
Department of Aeronautics and Astronautics
May 20, 2016

Signature redacted

Certified by .....
Hamsa Balakrishnan
Associate Professor
Thesis Supervisor

Signature redacted

Accepted by .....
Paulo C. Lozano
Associate Professor of Aeronautics and Astronautics
Chair, Graduate Program Committee

Analysis of Efficiency and Fairness in Stochastic Ground Holding Models

by

Karthik Kavassery Gopalakrishnan

Submitted to the Department of Aeronautics and Astronautics
on May 20, 2016, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

Allocating limited resources in a fair and efficient manner is crucial to the functioning of the air transportation system. Limited arrival and departure capacity at airports is one of the major drivers of congestion. The Ground Holding Problem refers to the problem of efficiently allocating landing slots during periods of reduced capacity. In this thesis, we propose a new model of the Ground Holding Problem. This formulation accounts for fairness in slot allocation and operational constraints, while being robust to the duration and severity of the disruption. The performance of this new method is evaluated using operational flight data.

Thesis Supervisor: Hamsa Balakrishnan
Title: Associate Professor

Acknowledgments

I thank my adviser Hamsa Balakrishnan for introducing me to exciting problems in the field of operations research and guiding me throughout this work.

It has been a pleasure interacting with Varun, who has provided numerous insights and suggestions while developing the model. Thank you for your time and advice.

Thanks to my friends in ICAT, especially Patrick, Jacob, Yash and Sandeep for creating a wonderful atmosphere in the lab.

Finally, I am glad to have a wonderful set of room mates and friends at MIT with whom I have shared countless adventures!

Contents

1 Introduction **13**

2 Ground Delay Programs **15**

2.1 Structure of a GDP 15

2.1.1 Ground Holding Problem 16

2.1.2 Collaborative Decision Making 18

2.2 Prior work on modeling GDPs 19

2.3 Scenario trees 21

2.4 GHP formulations 22

2.4.1 Static model 22

2.4.2 Dynamic model 24

2.4.3 Hybrid model 25

2.4.4 Example 26

2.5 Collaborative Decision Making (CDM) formulation 29

2.5.1 Intra-airline slot exchange 29

2.5.2 Slot Credit Substitution 31

3 Receding Horizon Static Ground Holding Problem **33**

3.1 RHS model 34

3.1.1 Terminology 35

3.1.2 Model 36

3.1.3 Number and time of updates 40

3.2 Intra-airline substitution under the RHS formulation 41

3.3 Integrated optimization formulation 43
3.4 Discussions 46

4 Results **49**
4.1 Preliminaries 49
4.1.1 Data set 49
4.1.2 GDP parameters 50
4.1.3 Probability distributions 51
4.2 Pre-CDM costs for existing models 53
4.3 RHS model 55
4.4 Intra-airline substitutions 56
4.5 Effect of cost variation 60
4.6 Slot credit substitution 62
4.7 Discussion 65

5 Conclusions **67**

List of Figures

2-1 Overview of a GDP. 16

2-2 Example scenario tree with a low capacity (10), and high capacity (20). The probability of occurence (p) of each of the five scenarios is shown on the right [25] 22

2-3 Scenario tree for comparison of the three existing GHP models [25]. 26

3-1 Example of resolved and unresolved scenarios. 36

4-1 ASPM data for LGA arrivals on Feb 17, 2014. 50

4-2 The $2T - 1$ probability distributions, with $T = 5$. 52

4-3 Static model costs for $T = 7$ (left), and $T = 15$ (right). 53

4-4 Hybrid model costs for $T = 7$ (left), and $T = 15$ (right). 54

4-5 Dynamic model costs for $T = 7$, (left) and $T = 15$ (right). 54

4-6 One update RHS for $T = 7$ (left), and $T = 15$ (right). 57

4-7 Two update RHS costs for $T = 7$ (left), and $T = 15$ (right) with $s = 1$. 58

4-8 Two update RHS costs for $T = 7$ (left), and $T = 15$ (right) with $s = T$. 58

4-9 Two update RHS costs for $T = 7$ (left), and $T = 15$ (right) with $s = 2T - 1$. 58

4-10 Cost of the RHS model in comparison to the static and dynamic model, for $T = 7$ (left), and $T = 15$ (right). 59

4-11 Optimal cost for different number of updates. Zero updates means that it is a static model. The different probability scenarios plotted are $s = 1$ (left), $s = 7$ (center), and $s = 13$ (right). $T = 7$ for all the three cases. 59

4-12 Post-CDM cost of all the four models for $T = 7$ (left), and the percentage improvement over the worst model (right). The missing bars correspond to the model with the highest cost. 61
4-13 Post-CDM cost of all the four models for $T = 15$ (left), and the percentage improvement over the worst model (right). The missing bars correspond to the model with the highest cost. 61
4-14 Actual costs (left), and percentage improvement over the worst model (right), for $T = 10$, coeff. of variation = 10%. 62
4-15 Actual costs (left), and percentage improvement over the worst model (right), for T=10, coeff. of variation = 30%. 62
4-16 Additional air holding (left), and ground holding (right) costs for the canceled flight. 64
4-17 Mean benefit for other flights (left), and for the system (right), when an SCS is triggered. 64

List of Tables

2.1 Example of a deterministic Ground Holding Problem. The initial (left), and revised schedule (right) is shown. 18

2.2 Example flight schedule. 26

2.3 Static slot reallocation. 27

2.4 Dynamic slot reallocation for Flight A. 27

2.5 Dynamic slot reallocation for Flight B. 27

2.6 Hypothetical case where Flight A takes the slot assigned to B, through a dynamic model. 28

2.7 Hybrid slot reallocation for Flight A. 28

2.8 Hybrid slot reallocation for Flight B. 28

2.9 Hypothetical case where Flight A takes the slot assigned to B, through a hybrid model. 29

3.1 Example of a RHS solution. 42

3.2 Comparison of the four GHP models. 47

4.1 Expected duration of reduced capacity for all the nine probability distributions. 51

Chapter 1

Introduction

Airports have a limited capacity in terms of how many aircraft can land or take off in a particular time interval. Estimating the capacity of the airport is non trivial because of several interacting factors [26, 16, 24]. The specific runway configuration being used, the weight of the landing aircraft, and the meteorological conditions (Visual or Instrument Flight Rules) all determine the minimum spacing required between aircraft. In addition, the landing and takeoff capacity of an airport is coupled because of the shared use of runways, taxiways and terminal regions. The above factors all constrain capacity. With growing air traffic, the number of operations is getting closer to this capacity limit.

During poor weather conditions, the capacity of the airport is further reduced. When such a capacity drop occurs, we have a demand-supply imbalance. A combination of several undesirable conditions may occur: departure queues getting built up, flights getting rerouted and the growth of airborne queues. In fact the effect of reduced capacity in a single airport can even cause delays at different airports because of passenger, crew and flight connectivity. Airlines, in coordination with the air traffic managers need to strategically respond in order to minimize disruptions.

Congestion management involves multiple time-scales. At the most fundamental level, demand forecasts (which occur upto a year in advance) helps airlines determine flight schedules. An estimate of capacity of the airport along with any existing slot restriction, yields the nominal airline flight schedules. Maximizing system and indi-

vidual airline efficiency are the primary drivers at this stage. On a more tactical level, which involves 2-10 hours into the future, a lot of new information becomes available. This includes details like weather forecasts, information regarding flight cancellations, mechanical faults and runway configurations. A revised demand and capacity estimate is generated. Traffic managers have several tools to manage congestion at this time-scale. On the airspace side, Traffic Flow Management programs may be put into place that restrict aircraft movements into an airspace. On the airport side, takeoffs destined for the congested airport may be restricted or completely stopped. These are known as Ground Delay Programs (GDP) and Ground Stops respectively. This thesis develops a methodology to model a GDP, accounting for fairness, efficiency and Collaborative Decision Making (CDM) among the stakeholders. Finally, at the shortest time scales of around 5-30 minutes, safety is the dominating factor in ATC decisions. Ensuring minimum separation and assigning a safe trajectory to each aircraft is the primary concern of Traffic Managers. They can hold arriving aircraft in the air, reroute them, divert them in case of emergencies, or assign different headings and altitudes to ensure separation.

To plan for a Ground Delay Program, traffic managers must first solve the Ground Holding Problem (GHP), which is a mathematical abstraction of the GDP. This is the focus of Chapter 2, in which the current methods to solve the GHP are described. The ideas behind CDM and its role in GDP planning is also illustrated. Chapter 3 describes a Receding Horizon Static (RHS) model which is the main contribution of this thesis. In Chapter 4, the proposed model is evaluated using operational data from LaGuardia airport in New York City, and comparisons are made with existing models. Chapter 5 summarizes RHS formulation, and also provides directions for further research.

Chapter 2

Ground Delay Programs

A GDP is initiated when the FAA's Enhanced Traffic Management System (ETMS) and Flight Schedule Monitor (FSM) projects a demand-capacity mismatch at a particular airport. When projected arrival traffic at an airport is greater than the number of aircraft that can land (Airport Arrival Rate or AAR), the Air Traffic Control System Command Center (ATCSCC) may issue a Ground Delay program. The aim is to reduce the number of arriving aircraft at the affected airport. Flights headed to the affected airport are given a revised departure time so that the aircraft arrival rate at the affected airport is reduced. This Traffic Management Initiative is called as a Ground Delay Program(GDP). In this chapter, we present an overview of GDPs, review the prior work in modeling GDPs, and discuss the importance of efficiency and fairness in GDP implementation.

2.1 Structure of a GDP

There are two steps in a Ground Delay Program (as described in Figure 2-1). They are:

1. Ground Holding Problem (GHP)
2. Collaborative Decision Making (CDM)

In the first step, the central planner (for instance, the FAA) uses flight schedules

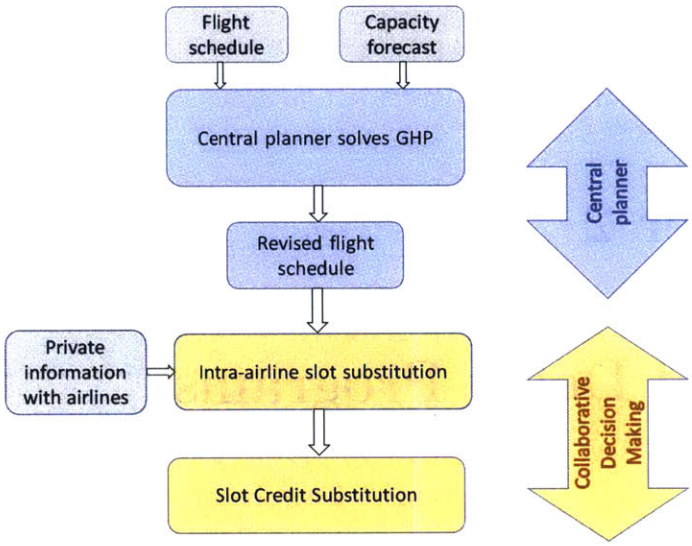

Figure 2-1: Overview of a GDP.

and capacity forecasts to solve the Ground Holding Problem and allocate an initial slot assignment to flights. In the second step (Collaborative Decision Making), airlines, who are stakeholders in the system, can modify the slot allocation to suit their needs. In this section, we will discuss these two steps in greater detail.

2.1.1 Ground Holding Problem

The Ground Holding Problem (GHP) can be described in the simplest form as follows: Given an anticipated decrease in airport capacity, how do we recompute a flight schedule that minimizes delay. The delays can be absorbed on the ground if the departure time is postponed. Or the delay can be absorbed in flight, for example, when the aircraft is put on hold in an airborne queue. If an aircraft has a revised departure time, the extra amount of time it spends at the airport before taking off is called the ground delay. When there are not enough available landing slots at an airport, an aircraft arriving to land has to wait for the next available slot in a holding pattern, which is usually a First In First Out (FIFO) queue. The amount

of time between landing at the airport and joining the airborne queue is called the airborne delay of a flight. Typically airborne delays are more expensive than ground delays because of fuel and maintenance costs. The ground holding problem takes into account the cost of air and ground holding in order to compute a minimum cost flight schedule.

GHPs can be formulated for two different cases, depending on how many airports are considered. The Multi Airport Ground Holding Problem deals with rescheduling flights when there is a decrease in capacity at several airports [34, 12, 35]. The problem is formulated as an Integer Program, and often relies on heuristics to deal with computational intractability [12]. In this thesis, we focus on the Single Airport Ground Holding Problem (SAGHP), where there is a decrease in capacity is expected at only at one airport. Henceforth, when we use the term GHP, we mean the SAGHP.

If the capacity prediction is deterministic, the problem is relatively straightforward. For example, consider the arrival schedule at an airport given in Table 2.1. Under normal operations, the capacity of the airport is 6 arrivals/hour, or an arrival every 10 minutes. Suppose the capacity from 7-8 pm is predicted to be 3 arrivals/hour because of a storm, and will be restored at 8 pm to 6 arrivals/hour. If no action is taken, flights $F2 - F6$ will have to wait in the air until a slot becomes available. The solution is to spread the flight arrivals such that the new arrival rate does not exceed capacity. The result is shown in Table 2.1. The initial order of flights is maintained and some flights are assigned ground delays. Maintaining the initial order of flights in the reassignment is considered to be fair by stakeholders. Later, we will generalize this notion of fairness to the stochastic problem. Also note that the deterministic solution in Table 2.1 is also a minimum cost solution (assuming that air hold costs are greater than ground hold costs).

However, capacity estimates are not deterministic. Weather forecasts and consequent arrival capacity estimates have a probability distribution that assign likelihood values to different capacities. This introduces a stochastic element to the problem.

A GDP issued by the ATCSCC for a particular airport will prescribe the following:

- Duration of the GDP: This is the duration for which arrivals into the airport

Flight Code	Scheduled Arrival	Rescheduled Arrival	Ground delay
F1	7:00	7:00	0 min
F2	7:10	7:20	10 min
F3	7:20	7:40	20 min
F4	7:30	8:00	30 min
F5	8:00	8:10	10 min
F6	8:10	8:20	10 min

Table 2.1: Example of a deterministic Ground Holding Problem. The initial (left), and revised schedule (right) is shown.

will be controlled. It can range from one hour to several hours.

- Scope of the GDP: Every GDP will have a scope, or the set of affected flights. It is common practice to exclude flights already in the air or international flights (except from Canada) from the GDP. The scope of a GDP is usually limited to airports within a certain radius of the issuing airport.

- Controlled Time of Departures: For flights within the scope of the GDP, a revised departure time is assigned. There are two possible approaches to doing this. The flight times may be assigned irrespective of which capacity scenario materializes or the flight time assignment may vary based on partial observation of realized capacities. These are referred as the *static* and *dynamic* formulation respectively. They are described in more detail later.

2.1.2 Collaborative Decision Making

In 1993, the FAA initiated the Federal Data Exchange project and this was the first step in the CDM approach. The idea behind CDM is to involve all the stakeholders in the decision making process. By sharing information and understanding the impact of decisions made at every level, it is expected that decisions will be closer to optimal. In the context of the ground delay program, CDM involves two processes - Intra-airline substitutions and Slot Credit Substitutions (SCS).

Intra-airline substitutions : The slots assigned to flights belong to the airline. This rule allows the airline to swap slots among its flights. Consider an airline with two of

its flights affected by a GDP and one flight being more delay sensitive than the other. There could be several reasons for the higher delay sensitivity. The flight could have a lot of passengers who will miss their connections in case of a delay. The aircraft could be larger and hence impact more passengers. This kind of information (value of individual flights) is proprietary to the airline and the CDM framework should allow the airline to accommodate varying flight priorities.

Slot Credit Substitution (SCS) : This aspect of CDM deals with flight cancellations and slot forfeiting. If an airline cancels its flight because it is unable to use the slot assigned to it, the schedule is compressed so that other aircraft can be pushed forward and the system benefits. In order for the airlines to be incentivized to reveal slot forfeitures or cancellations, the SCS mechanism ensures that the airline gets some benefit. A possible incentive that we discuss in this thesis is that an airline can request for any later slot it desires.

To summarize, we outline the process by which a GDP is issued and implemented. When the FAA predicts a demand supply imbalance at an airport, it issues a GDP at the airport. Based on the anticipated imbalance, the duration and score of the GDP is issued. A centralized optimization problem is solved involving all the flights in the scope. The result is an allocation of Controlled Departure Time (CDT) for the aircraft. The optimization is done such that the sum of the ground delay and airborne delay experienced by an airline is minimized. Each airline owns proprietary information regarding the relative value of each flight. Under the CDM paradigm, the airline can now swap slots internally. Further incentives are provided so that flight cancellations and slot forfeitures are reported.

In the next section, an overview of the literature on GDP's is presented.

2.2 Prior work on modeling GDPs

The deterministic ground holding problem is well studied [9, 30]. By deterministic, we mean that the estimates of future capacity are known exactly. Although this is a computationally tractable problem for a single airport case , it becomes intractable

when we do a multi-airport formulation [10], or include objectives such as minimizing passenger missed connections [5], and hubbing operations of airlines [18]. Heuristics have been proposed for the multi-airport GHP by Brunetta et al. and Vranas et al. [12, 34].

In the stochastic version of the ground holding problem, the airport capacities are described in the form of scenario trees. We describe scenario trees further in Section 2.3. In short, they assign a probability for each predicted AAR. Scenario free models [20] have not been widely implemented because of the computational difficulties. The earliest stochastic model was proposed by Richeta and Odoni [28]. It was a static model, wherein the GHP is solved and a Controlled Time of Departure (CTD) is assigned to every aircraft irrespective of how the capacity actually materializes. Thus, a fixed schedule is provided to the flights, that best accounts for the stochastic capacity, air hold and ground hold costs. In general, the static stochastic GHP formulation is not computationally tractable. However, when the ground holding costs are marginally non-decreasing, the LP relaxation has been shown to be integral [19]. Dynamic models, in contrast to static models, can make use of updated forecast information, and prescribe a more efficient ground hold. They do this by allocating a revised departure time for an aircraft at the latest possible moment, i.e., just before its scheduled departure. At this time, the maximum information regarding the state of the system would be available. Such a model was first proposed by Richeta and Odoni [29], and extended to the multi-airport case by Vranas et al. [35]. In the work by Richetta and Odoni [29], once a ground delay is assigned to a flight, it cannot be revised. Although this constraint results in more predictability, we can still replan the schedule of an aircraft until it takes off under the revised schedule. This observation was made by Mukherjee and Hansen [22], who also proposed an alternate dynamic formulation. The benefits of the Mukherjee and Hansen formulation over the Richetta Odoni model was shown to be the highest when either of the following occurred: Stringent ground holding, early cancellation of the GDP, or significantly reduced capacity. At this stage, a natural question that needs to be answered is: Why should the original static solution (that is computed before any scenario materi-

alizes), be followed even after new information is available? This question was partly answered by Mukherjee et al. [23] where the static solution was recomputed at every time interval in the GDP. This approach is called a Quasi Dynamic Stochastic Optimization (QDSO). There are several questions we need to still answer regarding such an approach. Is rerunning the static problem at every time instant the best strategy? What would happen if the static model was rerun a fewer number of times? How will these models integrate with the existing CDM paradigm? These questions motivate the development of the Receding Horizon Static (RHS) model that we will pursue.

Work has been done to address the challenges of fairness and equity in constrained resource allocation [6, 8, 7]. The centralized GHP problem has been reformulated such that the slot allocation can be subject to swaps or modifications by airlines. This could involve aggregate demand formulations [4], explicitly enforcing equality among flights of different duration [17], reducing exemption bias [32], incorporating a Ration By Distance formulation [2], or other rationing schemes [21]. From a implementaion perspective, the mediated bartering model [31, 33] and Top Trading Cycles Algorithm [1] address the intra-airline slot exchanges and compression step. Recent work by Cox and Kocenderfer [13, 14] proposes a CDM compliant Markov Decision Model. However the computational complexity of such an approach poses a challenge, even when simple cost functions are chosen.

2.3 Scenario trees

Scenario trees are one way to model stochastic capacity forecasts. Although scenario free methods for solving the GHP have been proposed [20], they present computational challenges. In this thesis, we use a similar structure of scenario trees as done by Ramanujam [27] (Figure 2-2).

The numbers in the circle denote the expected arrival capacity of the airport at that particular time. Every branch of the tree is a particular realization of the capacity. For example, the capacity estimate (10, 10, 20, 20, 20) occurs with a probability $\mathbb{P}(S3) = 0.4$. This means that the probability that the capacity improves

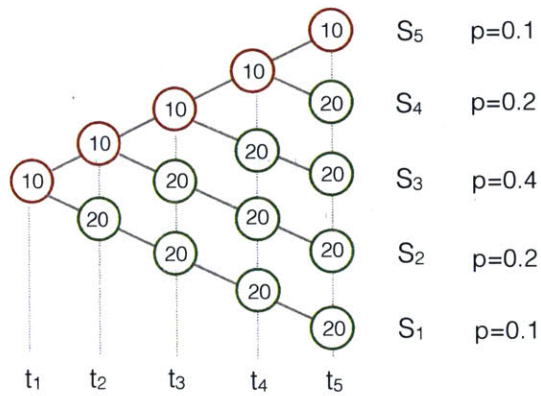

Figure 2-2: Example scenario tree with a low capacity (10), and high capacity (20). The probability of occurence (p) of each of the five scenarios is shown on the right [25]

at t_3 is 0.4. We assume that the time intervals between successive t_i's is 1 hour. In general they can be of any duration. The capacity at the end of a GDP is assumed to be large enough that any flight still delayed at the entire of the GDP can land immediately once the GDP is over (capacity at $t_6 \sim \infty$) .

Generating scenario trees may not be easy in practice. For example, we would require probabilistic forecasts of when a storm might clear. Also, translating a weather forecast to capacity estimates involves several factors such as intensity of the storm, runway configuration used and arrival-departure tradeoffs [25].

Now, we present the optimization formulations for the static, dynamic and hybrid models.

2.4 GHP formulations

2.4.1 Static model

The static GHP model with aggregate demands [3, 28] is presented.

Input data: The GHP is planned for time $t = 1, ..., T$. The stochastic capacity

forecast has $Q = \{1, ..., T\}$ scenarios. M_q^t is the capacity of the airport at time t under scenario q and π_q is the probability that scenario q materializes. N_t is the total number of aircraft that are scheduled to land at t. The average ground holding cost per aircraft for n time periods is $C_{g,n}$ and the average air holding cost per aircraft for a unit tine period is C_a. F is the set of all flights in the GHP.

Decision Variables: $X_{i,j}$ is the number of aircraft that were scheduled to land at time period i but are rescheduled to land at j. W_q^t is the number of aircraft that are airborne at time t in scenario q.

$$\min \quad \sum_{i=1}^{T} \sum_{j=i}^{T+1} C_{g,(j-i)} X_{i,j} + C_a \sum_{q=1}^{T} \pi_q \sum_{t=1}^{T} W_q^t \tag{2.1}$$

subject to

$$\sum_{j=i}^{T+1} X_{i,j} = N_i \quad \forall i = \{1, ..., T\} \tag{2.2}$$

$$W_q^t \geq \sum_{i=1}^{t} X_{i,t} + W_q^{t-1} - M_q^t \quad \forall t = \{1, .., T\}, q \in Q \tag{2.3}$$

$$W_q^0 = 0 \quad q \in Q \tag{2.4}$$

$$W_q^t \geq 0 \quad \forall t = \{1, .., T\}, q \in Q \tag{2.5}$$

$$X_{i,j} \geq 0 \quad \forall i = \{1, .., T\}, j = \{1, .., T+1\} \tag{2.6}$$

$$\tag{2.7}$$

The objective function (2.1) has two terms. The first is the deterministic cost that will come from the ground hold. The second term is the expected air holding cost. The airborne queue is scenario dependent, making it a random variable. Constraint (2.2) accounts for all aircraft in the GDP. Constraint (2.3) updates the airborne queue. When $C_{g,n}$ is a monotonically increasing function of n, we obtain a solution that follows the Ration By Schedule (RBS) principle.

2.4.2 Dynamic model

We describe the dynamic model of Mukherjee and Hansen [22]. Some additional variables are used in this model and they are explained below.

Notation: The arrival time for every flight $f \in F$ is arr_f and its flight duration is dur_f. G_t is the set of all scenarios still indistinguishable at time period t.

Decision variable : $X^q_{f,t} \in \{0,1\}$ is a binary decision variable. It is 1 when the flight f is rescheduled to land at time t under scenario q and 0 otherwise.

$$\min \quad \sum_{q=1}^{T} \pi_q (\sum_{f \in F} (\sum_{t=arr_f}^{T+1} C_{g,t-arr_f} X^q_{f,t}) + C_a \sum_{t=1}^{T} W^t_q) \tag{2.8}$$

such that

$$\sum_{t=arr_f}^{T+1} X^q_{f,t} = 1 \quad \forall q \in Q, \forall f \in F \tag{2.9}$$

$$W^t_q \geq W^{t-1}_q + \sum_{f \in F} X^q_{f,t} - M^t_q \quad \forall t = \{1,..,T\}, q \in Q \tag{2.10}$$

$$X^{q1}_{f,t} = X^{q2}_{f,t} \quad \forall q1, q2 \in G_{t-dur_f} \tag{2.11}$$

$$W^0_q = 0 \quad , \forall q \in Q \tag{2.12}$$

$$X^q_{f,t} \in \{0,1\} \tag{2.13}$$

$$W^t_q \in \mathbb{Z}^+ \tag{2.14}$$

By using flight specific decision variables that can assign revised schedules until the takeoff time ($arr_f - dur_f$), lower costs than the static model can be achieved. Constraint (2.9) requires that every scheduled flight gets a revised landing time, and Constraint (2.10) governs the queue dynamics for all scenarios. Decisions for any flight cannot be taken based on information not yet available (Constraint (2.11)). Constraints (2.12), (2.13) and (2.14) initialize the variables. The landing slots assigned by the dynamic model are dependent on flight-specific information, specifically the flight duration. We will see in Section 2.5 how this fact constrains the flight slot

swaps in the CDM step.

2.4.3 Hybrid model

The hybrid model described by Ramanujam and Balakrishnan [27] is now presented.

Notation: In addition to all the terms defined for the static and dynamic formulation, max_dur is the maximum flight duration among all flights that are considered for the GDP.

Decision Variable : $X^q_{i,j}$ is the number of flights that were scheduled to land at i, that are now rescheduled to land at j under scenario q.

$$\text{Minimize} \sum_{q \in Q} \pi_q \left(\sum_{i=1}^{T} \sum_{j=i}^{T+1} C_{g,j-i} X^q_{i,j} + C_a \sum_{t=1}^{T} W^t_q \right) \tag{2.15}$$

such that

$$\sum_{j=1}^{T+1} X^q_{i,j} = N_t \quad \forall i \in \{1,..,T\}, q \in Q \tag{2.16}$$

$$W^t_q \geq \sum_{i=1}^{t} X_{i,t} + W^{t-1}_q - M^t_q \quad \forall t \in \{1,..,T\}, q \in Q \tag{2.17}$$

$$X^{q1}_{f,t} = X^{q2}_{f,t} \quad \forall q_1, q_2 \in G_{t-max_dur} \tag{2.18}$$

$$W^0_q = 0 \quad \forall q \in Q \tag{2.19}$$

$$X^q_{i,j} \in \mathbb{Z}^+ \tag{2.20}$$

$$W^t_q \in \mathbb{Z}^+ \tag{2.21}$$

The use of aggregate decision variables, along with consideration of max_dur instead of dur_f, makes the formulation independent of individual flight schedules. This gives the hybrid formulation greater flexibility in slot swaps.

Flight Name	Departure Time	Arrival Time	Flight Duration
A	1	3	2
B	3	4	1

Table 2.2: Example flight schedule.

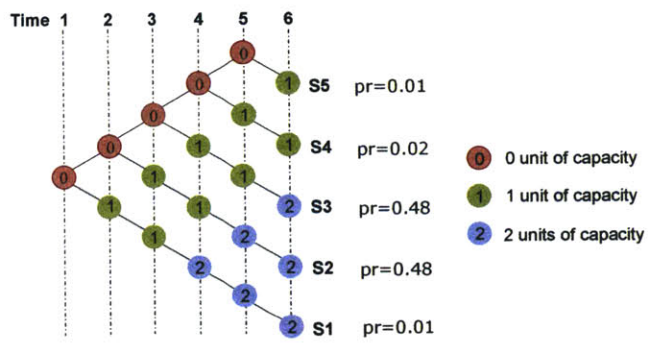

Figure 2-3: Scenario tree for comparison of the three existing GHP models [25].

2.4.4 Example

Ramanujam [25] describes the three models with the help of this example. Consider the scenario tree given in Figure 2-3 and the flight schedule of two aircraft described in Table 2.2.

If there was no ground holding policy, flight A would arrive at its destination at $t = 3$. With a probability 0.51, the capacity will be 0. When this happens, the flight will be put on an air hold until a landing slot is available. Suppose the ground holding cost be 1/aircraft/hour and air holding cost be 5/aircraft/hour. It is preferable to delay the aircraft on the ground than wait in the air. The static solution to this problem is presented in Table 2.3. Both aircraft are given a deterministic ground hold of 1 time period each. The aircraft arrive at $t = 4$ and $t = 5$ now. It is certain that there will be a ground hold of 1 time period for both the flights. However, there will be an air hold only in scenarios S4 and S5 (with probabilities 0.02 and 0.01 respectively).

The dynamic solution is presented in Table 2.4 and 2.5. Flight B gets different

Flight	Original Arrival Time	Ground Hold	Revised Arrival Time
A	3	1	4
B	4	1	5

Table 2.3: Static slot reallocation.

	S1	S2	S3	S4	S5
Original Departure	1	1	1	1	1
Original Arrival	3	3	3	3	3
Ground Hold	1	1	1	1	1
Revised Departure	2	2	2	2	2
Revised Arrival	4	4	4	4	4

Table 2.4: Dynamic slot reallocation for Flight A.

	S1	S2	S3	S4	S5
Original Departure	3	3	3	3	3
Original Arrival	4	4	4	4	4
Ground Hold	0	1	1	2	3
Revised Departure	3	4	4	5	6
Revised Arrival	4	5	5	6	7

Table 2.5: Dynamic slot reallocation for Flight B.

amount of ground holds depending on the scenario. The scheduled takeoff time for B is 3. At this time, S1 and S2 would have resolved themselves. If the scenario is S1, then B takes off on time. Else, it is ground held. At $t = 4$, if the scenario is S2 or S3, the flight takes off or else it is ground held. At $t = 5$, if updated information indicated the realization of S4, then B takes off, else it takes off at $t = 6$. Note that we assume a very high capacity at $t = 7$, so that any remaining flights affected by the GDP can land at that time.

Let us look at how an airline would swap slots under both the static and dynamic allocation. Suppose that Flights A and B are operated by the same airline, and that Flight B has a higher delay cost because it has a lot of connecting passengers. Swapping is straightforward in the static assignment. Flight A and Flight B are allotted a landing slot at $t = 4$ and $t = 5$ respectively. We will interchange these two slots. Flight A will land at $t = 5$ after incurring a ground delay of 2 units, and Flight

	S1	S2	S3	S4	S5
Original Departure	1	1	1	1	1
Original Arrival	3	3	3	3	3
Ground Hold	1	2	3	4	5
Revised Departure	2	3	3	4	5
Revised Arrival	4	5	5	6	7

Table 2.6: Hypothetical case where Flight A takes the slot assigned to B, through a dynamic model.

	S1	S2	S3	S4	S5
Ground Hold	1	1	1	1	1
Revised Departure	2	2	2	2	2
Revised Arrival	4	4	4	4	4

Table 2.7: Hybrid slot reallocation for Flight A.

	S1	S2	S3	S4	S5
Ground Hold	0	1	1	1	1
Revised Departure	3	4	4	4	4
Revised Arrival	4	5	5	5	5

Table 2.8: Hybrid slot reallocation for Flight B.

B will land at $t = 4$, incurring a ground delay of 0. Thus, Flight B takes off on time and the airline can attain its objective.

On the other hand, we will not be able to perform this slot swap for the dynamic allocation. This is because of the unequal flight durations associated with the two slots. If Flight A was to take the slot assigned to B, the schedule would look as shown in Table 2.6. This is an infeasible assignment since Flight A cannot distinguish between Scenario S4 and S5 at departure time 4. The slot was meant to be for a flight of duration 1, and not for a flight of duration 2. Thus swaps can only be made between flights of same duration.

The hybrid formulation addresses this limitation of the dynamic model. The resulting solution is presented in Table 2.7 and 2.8. If flight A were to take the slot for B, there would be no ambiguity in resolving any scenario (as seen in Table 2.9).

The optimal pre-CDM costs in the static, hybrid, and dynamic allocations in this

	S1	S2	S3	S4	S5
Ground Hold	0	1	1	1	1
Revised Departure	2	3	3	3	3
Revised Arrival	4	5	5	5	5

Table 2.9: Hypothetical case where Flight A takes the slot assigned to B, through a hybrid model.

example are found to be 2.4, 2.39, and 2.23, respectively. The pre-CDM cost is lowest for the dynamic model, followed by the hybrid and static model.

2.5 Collaborative Decision Making (CDM) formulation

The response of the airlines to the initial slot allocation is discussed. This was first presented by Ramanujam (2011) [25]. We formulate the intra-airline slot substitution as an assignment problem. Then we discuss the incentives required for truthful reporting of slot cancellations.

2.5.1 Intra-airline slot exchange

Once a GHP is solved by the FAA (using nominal values for C_a and $C_{g,n}$), the following information is sent to airlines for each of their flights

1. Minimum ground holding time: The minimum ground hold that will be assigned to a slot s across all scenarios is $grd_{min}(s)$.

2. Expected ground holding time: This is the average ground holding time that will be assigned to the slot. For a slot s, the expected time is $grd(s)$.

3. Expected air holding time: The airborne queue that a slot will experience is dependent on the scenario. The scenario probability weighted mean wait time for a slot s is denoted as $air(s)$.

The airlines have flight-specific delay costs. For each flight f, the air delay cost for each time period is C_a^f, and the ground hold cost per time periods is C_g^f. In the original schedule, every flight also has a departure time $dep(f)$, along with a flight duration $dur(f)$. The aim is to find a minimum cost matching of flights to slots.

Let $\mathcal{F}_a$ be the set of flights assigned to airline a. Denote the flight corresponding to slot k as f_k. The cost of assigning a flight f to slot associated with flight k (with an associated lot s_k) is given by

$$C_{f,k} = C_g^f grd(s_k) + C_a^f air(s_k) \quad f, f_k \in \mathcal{F}_a \tag{2.22}$$

We define a 0-1 indicator variable to determine the feasibility of a slot swap. $feas(f,k)$ is 1 if the flight f can be allocated to slot s_k. A necessary condition for feasibility is that the flight cannot take off before its original published time. This is used as a surrogate for the earliest possible departure time of the flight. The necessary condition for feasibility of swaps ($feas(f,k) = 1$), is $dep(f) \geq dep(f_k) + grd_{min}(s_k)$. For the static and hybrid formulation, this is also a sufficient condition. For the dynamic formulation, we have an additional restriction on the flight duration, i.e., $dur(f) = dur(f_k)$.

The optimization problem finds the minimum cost, scenario-independent assignment of flights to slots.

$$\min \sum_{f \in \mathcal{F}_a} \sum_{k: f_k \in \mathcal{F}_a} C_{f,k} \tag{2.23}$$

subject to

$$\sum_{f \in \mathcal{F}_a} X_{f,k} = 1 \quad \forall \{k : f_k \in \mathcal{F}_a\} \tag{2.24}$$

$$\sum_{k: f_k \in \mathcal{F}_a} X_{f,k} = 1 \quad \forall f \in \mathcal{F}_a \tag{2.25}$$

$$X_{f,k} \leq feas(f,k) \quad \forall f \in \mathcal{F}_a, \{k : f_k \in \mathcal{F}_a\} \tag{2.26}$$

$$X_{f,k} \in \{0,1\} \quad \forall f \in \mathcal{F}_a, \{k : f_k \in \mathcal{F}_a\} \tag{2.27}$$

Equations (2.24) and (2.25) are required so that a one-one mapping between flights and slots is maintained. Equation (2.26) ensures only feasible swaps.

2.5.2 Slot Credit Substitution

Airlines might also respond to GDPs by canceling flights. Sometimes, it may be cheaper to cancel a flight instead of delaying it by a lot. Since a flight has multiple legs in its journey, it is also possible that an aircraft might not be able to use a slot assigned to it. The Slot Credit Substitution (SCS) mechanism is triggered whenever an airline decides to forfeit a slot. This is in contrast to the earlier compression algorithms that were run in batches.

Ramanujam [25] describes three essential features of a SCS mechanism in the stochastic context:

1. An airline which forfeits a slot c can request any later slot k.

2. All other flights see a decrease in air and ground delay (Pareto improvement for all other flights).

3. The benefit of a cancellations must be distributed to all the flights in a equitable manner.

The first two conditions ensure an incentive compatible mechanism for the airlines to report slot cancellations, since they are going to be strictly better off if they report a cancellation, or request a later slot. The last feature refers to the equitable manner in which the benefits of a slot forfeiture are distributed among other airlines. Pareto improvement requires that each flight has no increase in both the airborne queue and the ground hold. This is a strong requirement. Although not addressed in this thesis, it is possible to develop a broad definition of "improvement" in terms of airline benefits. In the simulations described in Chapter 4, we will compute the system wide benefits of implementing the SCS algorithm described in [25].

For the SCS optimization formulation in [25], the input data is the following:

- The flight that is forfeiting its current slot assignment, c, and the later time slot that it requests, k
- Information regarding all the other flights
  - $ETA(slot_f)$: The earliest arrival time of each aircraft assigned to slot f
  - W_q^t: The airborne queue length
  - $dur(slot_f)$: Duration of the slot assigned to flight f
  - $arr^q(slot_f)$: Arrival time of flight f under scenario q

The SCS optimization will try to allot a slot at time k or later to the airline that vacates c, and decrease the delay for all other feasible flights. The benefit (decrease in delay) for the canceled flight, the non-canceled flights, and the entire system, are computed in Chapter 4.

Chapter 3

Receding Horizon Static Ground Holding Problem

In this chapter, we introduce the Receding Horizon Static (RHS) formulation for the stochastic SAGHP. In the previous chapters, we studied the features of the static, dynamic and hybrid models. The static model was found to provide the least pre-CDM benefits. However, it is the most flexible in terms of slot exchanges, and gives the most benefit during the CDM step. By contrast, the dynamic model has the best pre-CDM performance because it takes advantage of the latest capacity forecasts. But the flight duration-specific slot assignment constrains slot swaps, resulting in the least CDM benefits. This tradeoff motivated the development of the hybrid formulation [27].

The hybrid formulation can make use of dynamic information available at every time step and also ensure that slot substitutions are not restricted. Like the dynamic model, there is a large communication overhead associated with it. At every time step, the capacity estimates need to be shared among all the stakeholders. The current branch of the scenario tree must also be common knowledge. Operationally, this can pose a challenge. Airlines cannot plan for future time steps because of the uncertainties involved.

We list the features that we would like to see in an ideal GDP algorithm:

- Minimum cost: The primary aim of the optimization is to try and transfer all the delay to the ground instead of the air, and balance it against the risk of starving the arrival runway. The Post-CDM cost is the primary metric for evaluation (as opposed to the pre-CDM cost)

- Fair allocation of resources: This is an important feature for airlines. The ideal allocation will not distinguish between flights of different duration, or airlines with different market shares. Ration-By-Schedule (RBS) is accepted as a fair allocation procedure by the airline industry.

- Robustness of algorithm: The length, severity (reduced capacity to nominal capacity ratio), and scope of GDPs vary widely. The algorithm must therefore show consistent performance in a wide range of scenarios

- Predictability: Airlines would like to know the departure time of their flights well in advance. As per this consideration, static models are most preferable. Also, once assigned, the schedule should preferably not change (which supports the dynamic formulation of Richetta and Odoni [29] over that of Mukherjee and Hansen [22]).

- Computational tractability of the formulation: The problem should not scale exponentially with the number of flights involved.

3.1 RHS model

The RHS model is motivated by the following two favorable features of existing models:

- The static model gives a slot assignment well in advance of the GDP, and also provides maximum flexibility to make intra-airline slot substitutions. Further, the static solution will also be a fair RBS allocation if the ground holding costs are marginally non-decreasing [19].

- The strength of the dynamic formulation comes from its ability to use updated capacity forecasts. We incorporate this feature into the RHS model by allowing for multiple runs of the static model. Whenever a static solution is obtained, it is implemented by airlines until the next run of the static optimization.

In this process, we hope to develop an alternative to the hybrid model which tries to achieve the same purpose: a low post-CDM cost.

3.1.1 Terminology

Let the length of the GDP be T. The discrete time-steps are $\{1, 2, ..., T\}$. The static optimization is re-solved with updated information at specific times, denoted as the "update times". Therefore, if we have k update times, t_1, .., t_k, with $1 < t_1 < ... < t_k < T$, then the static model will be run at times 1, t_1, t_2,..., t_k. Whenever the static model is run at time t_i, a schedule is planned for times t_i through T. The period between two updates is denoted as a stage. The first stage is therefore from $t = 1$ to $t = t_1 - 1$, the second stage from $t = t_1$ to $t = t_2 - 1$, and the $(k+1)^{th}$ stage is from $t = t_k$ to $t = T$.

Consider the scenario tree in Figure 3-1. Let the update times be $t_1 = 2$ and $t_2 = 4$. At $t_1 = 2$, the occurrence or non-occurrence of S1 is known. Scenario S1 is said to be resolved at time t=2. Scenarios S2 to S6 are still indistinguishable and are said to be unresolved at that time. Now when the second update happens at $t_2 = 4$, S2 and S3 will get resolved (in addition to S1, which was already resolved). At t_2, scenarios S4 to S6 remain unresolved. In general, at any time t, scenarios S1,...,S(t-1) would have been resolved.

We define a k-stage RHS formulation as one where the static model is run k times. Since the model is always run at $t = 1$, there will be $k - 1$ update times. Specifically, a 2-stage RHS model has one update time.

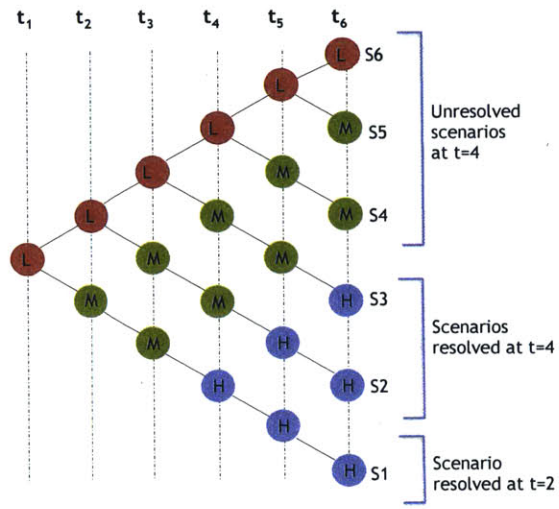

Figure 3-1: Example of resolved and unresolved scenarios.

3.1.2 Model

The RHS formulation is similar to the static model in many aspects. We however do not use the decision variable $X_{i,j}$ as shown in Section 2.4.1. In RHS implementation, we need to keep track of flights across stages, and it is therefore important to know the flight duration. We will instead use the variable $X_{s,i,j}$ (defined below). It is to be noted that this decision variable is not flight-specific. The formulation will not distinguish between flights of different durations, unlike the dynamic model. This variable is just needed for book-keeping purposes.

The RHS model is a generic framework based on rerunning the static model. However, we will only present the details for a 2-stage version. The basic ideas developed here can be generalized for any number of updates.

Step 1: Solve the static model at t=1

$$\text{Minimize} \quad \sum_{s=1}^{T}\sum_{i=1}^{T}\sum_{j=i}^{T+1} C_{g,j-i}X^{I}_{s,i,j} + \sum_{q\in Q}\pi_q(C_a\sum_{t=1}^{T}W_t^{\mathrm{I,q}}) \tag{3.1}$$

$$\text{subject to} \quad \sum_{j=i}^{T+1} X^{I}_{s,i,j} = N_{s,i},\ \forall s \in \{1,...,T\}, \forall i \in \{1,..,T\} \tag{3.2}$$

$$W_t^{I,q} \geq \sum_{s=1}^{T}\sum_{i=1}^{t} X^{I}_{s,i,t} + W_{t-1}^{I,q} - M_t^q,\ \forall t \in \{1,..,T\}, q \in Q \tag{3.3}$$

$$W_0^{I,q} = 0,\ \forall q \in Q \tag{3.4}$$

$$X^{I}_{s,i,j} \in \mathbb{Z}^+,\ W_t^{I,q} \in \mathbb{Z}^+,\ \forall t,s,i,j \in \{1,..,T\},\ \forall q \in Q \tag{3.5}$$

Notation:

$X^{I}_{s,i,j}$: Number of flights scheduled to take off at s and land at i, which are rescheduled to land at j. This is the variable that contains the stage I solution of the GDP (decision variable)

$W_t^{I,q}$: Airborne queue at time t for scenario q (decision variable for stage I of the GDP)

$\tilde{t}$: Update time for the 2-stage RHS model

$N_{s,i}$: Number of flights scheduled to takeoff at time s and land at time i (given from the original schedule)

Q : Set of all scenarios $\{1,...,T\}$

Equations (3.1)-(3.5) are similar to the static formulation in Section 2.4.1.

Step 2: Implement the solution until $\tilde{t}-1$

$X^{I}_{s,i,j}$ and $W_t^{I,q}$ is the solution to the first optimization. All flights with a rescheduled take off time before $\tilde{t}$ will follow the static solution. At $t=\tilde{t}$, the optimization is going to be rerun.

Step 3: Run the second static optimization

The second run of the static model needs to consider the following issues at $\tilde{t}$:

1. *Flights that have been ground held in stage I and have not yet departed*: A flight ground held through the first stage is reconsidered in the second stage optimization problem. They are considered as new flights scheduled to take

off at $\tilde{t}$. By revising their take off time to $\tilde{t}$, these flights will get preference over the other second stage flights and we retain the RBS property of the static formulation. It will get a revised takeoff time depending on capacity estimates. Define an auxiliary flight schedule matrix for stage II as $N^{II}_{s,i}$.

$$N^{II}_{s,i} = N_{s,i} \quad \forall s, i \geq \tilde{t} \tag{3.6}$$

$$N^{II}_{\tilde{t},i-s} = N^{II}_{\tilde{t},i-s} + X^{I}_{s,i,j} \quad \forall s < \tilde{t}, (s+j-i) \geq \tilde{t} \tag{3.7}$$

Equation (3.6) considers the original schedule and (3.7) accounts for flights which are ground held.

2. *Flights that have taken off at stage I but are currently airborne*: Let the variable $R_t \in \mathbb{Z}^+, t > \tilde{t}$ denote the number of airborne flights from stage I which will be landing in stage II. It is calculated according to (3.8)

$$R_t = R_t + X_{s,i,t} \quad , t > \tilde{t}, \{(s,i) : (s+t-i) < \tilde{t} < t\} \tag{3.8}$$

3. *The airborne queue from the first stage which gets handed over to the second stage solver*: We add the constraint that the initial queue for the second stage is equal to that obtained form the first stage solution.

4. *Scenario tree update*: The second stage optimization ((3.11)-(3.15)) involves two possibilities: a specific scenario is realized, or scenarios are unresolved. When a specific scenario, say q^* gets resolved, then $\tilde{Q} = \{q^*\}$ and $\tilde{\pi}_{q^*} = 1$ (because it reduces to a deterministic problem). The probability of a scenario being unresolved at $t = \tilde{t}$ is denoted by p_{unres} and in this case, we have $\tilde{Q} = \{\tilde{t}, ..., T\}$. The probabilities being used in the formulation are then conditionally updated as described in (3.9)-(3.10).

$$p_{unres} = \sum_{q=\tilde{t}}^{T} \pi_q \tag{3.9}$$

$$\tilde{\pi}_q = \frac{\pi_q}{p_{unres}} \quad \forall q \in \{\tilde{t}, ..., T\} \tag{3.10}$$

5. *Total cost of the two stage solution*: This is evaluated after we have both the stage solutions.

The second stage of the static optimization is solved considering the updated inputs. The main difference in the second stage formulation are the use of the updated flight schedule (Constraint (3.12)) and explicit accounting of airborne flights (Constraint (3.13)).

$$\text{Minimize} \quad \sum_{s=\tilde{t}}^{T}\sum_{i=\tilde{t}}^{T}\sum_{j=\tilde{t}}^{T} C_{g,j-i} X^{II}_{s,i,j} + \sum_{q\in\tilde{Q}} \tilde{\pi}_q (C_a \sum_{t=\tilde{t}}^{T} W^{II,q}_t) \tag{3.11}$$

$$\text{subject to} \quad \sum_{j=i}^{T} X^{II}_{s,i,j} = N^{II}_{s,i}, \ \forall s, i \geq \tilde{t} \tag{3.12}$$

$$W^{II,q}_t \geq \sum_{s\geq\tilde{t}}\sum_{i=\tilde{t}}^{t} X^{II}_{s,i,t} + W^{II,q}_{t-1} - M^q_t + R_t, \ \forall t > \tilde{t}, q \in \tilde{Q} \tag{3.13}$$

$$W^{II,q}_{\tilde{t}} = W^{*I,q}_{\tilde{t}}, \ \forall q \in \tilde{Q} \tag{3.14}$$

$$X^{II}_{s,i,j} \in \mathbb{Z}^+, \ W^{II,q}_t \in \mathbb{Z}^+, \ \forall t, s, i, j \geq \tilde{t}, \ \forall q \in \tilde{Q}, \tag{3.15}$$

Notation:

$X^{II}_{s,i,j}$: Flights scheduled to take off at s and land at i, which are rescheduled to land at j. This solution is in stage II of the GDP (decision variable).

$W^{II,q}_t$: Airborne queue for stage II of the GDP planning at time t for scenario q (decision variable).

$N^{II}_{s,i}$: Number of flights scheduled to takeoff at time s and land at time i in the second stage (from Equation (3.6) and (3.7)).

$\tilde{Q}$: Denotes the set of scenarios that are considered for the second stage.

$\tilde{\pi}_q$: Probabilities associated with the resolved or unresolved scenarios in set $\tilde{Q}$.

$W^{*I,q}_{\tilde{t}}$: Optimal solution of the first stage airborne queue in scenario q.

The total delay cost for a 2-stage RHS is computed as the sum of the Stage I (given in Equation (3.16)) and Stage II costs. The optimal solutions from the two

stages are $X^{*,I}_{s,i,j}$, $X^{*,II}_{s,i,j}$, $W^{*I,q}_t$, and $W^{*II,q}_t$. The stage I cost is the ground holding costs from $t = 1$ to $t = \tilde{t} - 1$ and the expected air delay for the same time period.

$$Cost_I = \sum_{s<\tilde{t}} \sum_{i>s} \sum_{j\geq i} C_{g,min(j-i,\tilde{t}-s)} X^{*,I}_{s,i,j} + \sum_{q\in Q} \pi_q (C_a \sum_{t=1}^{\tilde{t}-1} W^{*I,q}_t) \tag{3.16}$$

The Stage II cost has two cases, one in which the scenario is resolved, and another in which the scenarios are unresolved. When a particular scenario q is resolved in Stage II, the cost is given by:

$$Cost_{II,q} = \sum_{s\geq\tilde{t}} \sum_{i>s} \sum_{j\geq i} C_{g,j-i} X^{*,II}_{s,i,j} + C_a \sum_{t\geq\tilde{t}} W^{*II,q}_t \quad \forall \ resolved \ q \tag{3.17}$$

When the scenarios are not resolved, we update the expected cost as:

$$Cost_{II,unresolved} = \sum_{s\geq\tilde{t}} \sum_{i>s} \sum_{j\geq i} C_{g,j-i} X^{*,II}_{s,i,j} + \sum_{q\in\tilde{Q}} \tilde{\pi}_q (C_a \sum_{t\geq\tilde{t}} W^{*II,q}_t) \tag{3.18}$$

The total cost of stage II is therefore a probability-weighted sum of the resolved and unresolved scenarios:

$$Cost_{II} = p_{unres} * Cost_{II,unresolved} + \sum_{q=1}^{\tilde{t}-1} \pi_q * Cost_{II,q} \tag{3.19}$$

The total cost of the solution is:

$$Total \ Cost = Cost_I + Cost_{II} \tag{3.20}$$

3.1.3 Number and time of updates

An important question regarding the RHS model is the choice of the update time. If the update time is $t = 1$, then we gain no additional information about scenario occurrences. Further, if the update time is $t = T$, then all flights have already followed the schedule decided by the first stage. Despite of the ability to identify the scenario precisely, there are no flights that can benefit. There is a tradeoff between

early updates, with which we can potentially control a lot of flights in stage II, and waiting for better capacity estimates. This optimal update time depends on the flight schedules, reduced capacity estimates, probability distribution among the scenarios, and the length of the GDP. We do not study theoretical estimates, or formulate this as an optimization problem in this thesis. Instead, we identify the optimal update times through simulations. We illustrate this process with an example in the next chapter.

The other important parameter is the number of updates. We described the 2-stage RHS model. If we increase the number of stages, we obtain better capacity forecasts, which will decrease the pre-CDM costs. Intuitively, if we have a static update at every time step, we will have the lowest cost. However, this solution still has a higher cost than the dynamic solution. When we have multiple updates to make, we need to search from all possible permutations to find the optimal update times. For example, if we decide to have k updates for a GDP of time horizon T, we must find the update times from $\frac{T!}{k!(T-k)!}$ combinations. To summarize, the pre-CDM costs have the following order

$$Cost[dynamic\ GHP] \leq Cost[T{-}updates] \leq \ldots \leq Cost[2{-}updates] \leq Cost[static\ GHP]$$

3.2 Intra-airline substitution under the RHS formulation

We will first summarize how the 2-stage RHS formulation is implemented. Suppose a GDP of a certain duration and scope has been called due to a predicted supply-demand imbalance. Since the system operator does not have airline-specific costs for ground and air delays, for the sake of fairness, it assumes the same value for all flights. It then solves a 2-stage RHS model, and determines the minimum delay cost solution and the optimal update time. The system operator has the solution for the first stage, and a scenario-specific solution for the second stage. The following information is now conveyed to the airlines:

Flight	Departure time	Duration	Arrival Time
F1	5 or 6	1	6 or 7
F2	2	5	7

Table 3.1: Example of a RHS solution.

- The update time $\tilde{t}$.
- The flights which are scheduled to take off before $\tilde{t}$ and their actual takeoff times.
- The expected and earliest takeoff times for flights in the second stage.
- The expected air holding times for all flights.

Given this assignment and information on delays, flight swaps are done prior to the start of the GDP itself. We restrict ourselves to the case where airlines make only one round of intra-airline swaps. Technically, the swap algorithm can be rerun every time we use updated information but it is not practical from an operational perspective. An important difference between the RHS and the static model is the feasibility of swaps across stages. If two flights are taking off in the same stage, then we can do a swap independent of the flight duration. This is not possible across stages, in spite of both being static solutions.

Let us look at an example which illustrates this. Let the scenario tree structure be as before i.e at time t, scenarios 1,..., $t-1$ have been resolved (Figure 2-2). There are two flights, F1 and F2, with their RHS solution given in Table 3.1. The update time is $\tilde{t} = 4$, so flight F1 can take off at either $t = 5$ or at $t = 6$, depending on the scenarios that are resolved at $t = 4$. Suppose F2 is a delay-sensitive flight and we want to decrease its delay, and let it land earlier at $t = 6$ with a non-zero probability. If F1 and F2 swap slots, there is a conflict. F2 must now depart at $t = 1$ or $t = 2$ depending on a scenario that will only get resolved at $\tilde{t} = 4$. Hence the swap is infeasible.

A flight f can feasibly take another slot assigned to flight k if either of the following conditions are met:

1. Flights f and k are both rescheduled to take off in the first stage and the scheduled departure of flight f is before the earliest departure time of slot s_k, i.e. $dep(f) \geq dep_{min}(s_k)$

2. Fights f and k are both rescheduled to take off in the second stage and the departure time of f is before the earliest slot time for s_k, i.e. $max\{dep(f), \tilde{t}\} \leq dep_{min}(s_k)$. There is a max operator so that second stage flights do not take off before $\tilde{t}$.

These conditions reinforce the inherent trade-off between getting an optimal solution (by using all available information in a dynamic setting) and the flexibility offered in swapping slots.

3.3 Integrated optimization formulation

The RHS formulation, as described in the previous section, does not account for updates in its planning. In other words, the first stage of the static solver does not know that it will not be implemented fully, and that a second solution will be obtained at $\tilde{t}$. This feature is what fundamentally distinguishes even the T-stage RHS from a dynamic model. In this section, we present an integrated formulation for a 2-stage RHS.

The optimal solution for the original RHS model is a feasible solution to the integrated formulation as well. The integrated formulation will have a lower cost than the original RHS model. The drawback of the integrated formulation is that delay distribution among flights may not be fair. The first stage, accounting for the second update will tend to delay short haul flights when approaching the update time. Duration-based slots will reduce flexibility for slot swaps, and that will lead to lower CDM benefits.

A single optimization problem is set up for a known update time, $\tilde{t}$. The input parameters for the optimization are T, $N_{s,i}$, M_q^t, π_q, $C_{g,n}$ and C_a.

The decision variables in this formulation are:

$X^1_{s,i,j}$: Number of flights scheduled to take off at s and land at i, which are now rescheduled to land at j during the first stage.

$W^1_{q,t}$: The airborne queue at time t in scenario q during the first stage.

$X^2_{s,i,j}$: Number of flights scheduled to take off at s and land at i, which are now rescheduled to land at j during the second stage of unresolved scenarios.

$W^2_{q,t}$: The airborne queue at time t in scenario q during the second stage of unresolved scenarios.

$X^{(q)}_{s,i,j}$: Number of flights scheduled to take off at s and land at i, which are now rescheduled to land at j when solving for the qth resolved stage.

$W^{(q)}_t$: The airborne queue at time t in the qth resolved stage.

r_i : The number of airborne flights at time t^* which are scheduled to arrive at time i.

max_dur : It is the maximum duration of any flight in the GDP.

When the update time is $\tilde{t}$, the scenarios $Q_{res} = \{1, 2, ..., \tilde{t} - 1\}$ are resolved and $Q_{unres} = \{\tilde{t}, \tilde{t} + 1, ..., T\}$ are unresolved. $Q = \{1, 2, ..., T\}$ is the set of all possible scenarios. Also the total number of scenarios $\|Q\| = T$.

The objective function has three cost components coming from the first stage, the unresolved scenarios in the second stage and the resolved scenarios. We define

$$A = \sum_{s<t^*} \sum_{i=s+1}^{T} \sum_{j=i}^{T+1} C_{g,min\{(j-i),(t^*-s)\}} X^1_{s,i,j} + C_a \sum_{q=1}^{Q} \pi_q \sum_{t=1}^{t^*-1} W^1_{q,t} \tag{3.21}$$

$$B = \sum_{q=1}^{t^*-1} \pi_q \left[\sum_{s \geq t^*} \sum_{i=s+1}^{T+1} \sum_{j=i}^{T+1} C_{g,(j-i)} X^{(q)}_{s,i,j} + C_a \sum_{t=t^*}^{T} W^{(q)}_t \right] \tag{3.22}$$

$$C = \sum_{k=t^*}^{T} \pi_k \left[\sum_{s \geq t^*} \sum_{i=s+1}^{T+1} \sum_{j=i}^{T+1} C_{g,(j-i)} X^2_{s,i,j} + C_a \sum_{q=t^*}^{T} \frac{\pi_q}{\sum_{i=t^*}^{T} \pi_i} \sum_{t=t^*}^{T} W^2_{q,t} \right] \tag{3.23}$$

A is the cost of the first stage. The two terms correspond to the ground holding cost and the air holding cost respectively. B is the scenario specific second stage cost, and C is the cost of the unresolved scenarios. This objective is similar to the total cost expression in (3.16)- (3.20). The objective is to minimize $A + B + C$.

We describe the constraints for the dynamics of the first static stage. Equations (3.24)-(3.26) are the usual static equations that have the schedule, the queue length

evolution for different scenarios, and the initial conditions.

$$\sum_{j=i}^{T+1} X^1_{s,i,j} = N_{s,i} \quad \forall s < t^*, i = \{s+1, ..., T\} \tag{3.24}$$

$$W^1_{q,t} \geq W^1_{q,t-1} + \sum_{s<t} \sum_{i=(s+1)}^{t} X^1_{s,i,t} - M_{q,t} \quad \forall q \in Q, t = \{1, 2, .., t^*\} \tag{3.25}$$

$$W^1_{q,0} = 0 \quad \forall q \in Q \tag{3.26}$$

The next set of constraints link this first stage solution to the second stage. There is a continuity in the queue length (Equation (3.27)) across stages, an account of the flights pushed from first stage to the second (Equations (3.28) and (3.29)), and a variable r_t that tracks airborne flights (Equation (3.30)).

$$W^2_{q,t^*} = W^1_{q,t^*} \quad \forall q \in Q \tag{3.27}$$

$$N^2_{s,i} = N_{s,i} \quad \forall s > t^*, i > t^* \tag{3.28}$$

$$N^2_{t^*,t^*+n} = N_{t^*,t^*+n} + \sum_{s<t^*} \sum_{j=t^*+1}^{T+1} X^1_{s,s+n,j} \quad \forall n \in \{1, ..., max_dur\} \tag{3.29}$$

$$r_t = \sum_{(s,i):s<t^*, s+(t-i)<t^*} X_{s,i,t} \quad \forall t \in \{t^*+1, ..., T+1\} \tag{3.30}$$

The second stage constraints depend on whether the scenarios are resolved or not. The set of constraints for the deterministic scenarios are given by Equations (3.31)-(3.33) .

$$\sum_{j=i}^{T+1} X^{(q)}_{s,i,j} = N^2_{s,i} \quad \forall q \in Q_{res}, s \geq t^*, i > s \tag{3.31}$$

$$W^{(q)}_t \geq W^{(q)}_{t-1} + \sum_{s=t^*}^{t-1} \sum_{i=s+1}^{t} X^{(q)}_{s,i,t} - M_{q,t} + r_t \quad \forall q \in Q_{res}, t > t^* \tag{3.32}$$

$$W^{(q)}_{t^*} = W^1_{q,t^*} \quad \forall q \in Q_{res} \tag{3.33}$$

The unresolved scenarios have the following constraints:

$$\sum_{j=i}^{T+1} X^2_{s,i,j} = N^2_{s,i} \quad \forall s \geq t^*, i \in \{t^*+1, ..., T+1\} \tag{3.34}$$

$$W^2_{q,t} \geq W^2_{q,t-1} + \sum_{s=t^*}^{t-1} \sum_{i=s+1}^{t} X^2_{s,i,t} - M_{q,t} + r_t \quad \forall q \in Q_{unres}, t > t^* \tag{3.35}$$

$$W^2_{q,t^*} = W^1_{q,t^*} \quad , \forall q \in Q_{unres} \tag{3.36}$$

Finally, we require that the variables are all non-negative integers.

$$X^1_{s,i,j}, X^2_{s,i,j}, X^{(q)}_{s,i,j}, W^{(q)}_t, W^1_{q,t}, W^2_{q,t}, r_t \quad \in \mathrm{Z}^+ \tag{3.37}$$

The optimization problem is to maximize the objective $(A + B + C)$ subject to Constraints (3.24)-(3.37).

3.4 Discussions

The hypothesized performances of the four algorithms are compared in Table 3.2. The RHS model offers a simpler alternative to the hybrid formulation. However, it is not trivial to compare the hybrid and RHS models theoretically. The RHS solution is not even feasible to the hybrid optimization, because of scenario-specific decisions that are taken at the update time. Also the Hybrid solution is infeasible when plugged into the RHS optimization since it allows scenario based slots (if done max_dur time steps before). In the next chapter, we present pre- and post-CDM comparisons of all four models. The performance of the RHS and Hybrid formulation will also be discussed.

Model	Pre-CDM cost	CDM Benefits
Static	High	High
Hybrid	Medium	Medium
RHS	Medium	Medium
Dynamic	High	Low

Table 3.2: Comparison of the four GHP models.

Chapter 4

Results

In this chapter, we simulate GDPs to illustrate the different solution approaches, and their outcomes. First, we present the pre-CDM delay costs for all the models across a range of GDP conditions. The effect of choosing an optimal update time for the RHS model is highlighted. Next, intra-airline slot substitution is performed, and the post-swap costs are plotted. Finally, the slot credit substitution (SCS) process is explained in greater detail, and the benefits of the SCS procedure is compared across the different methods.

4.1 Preliminaries

4.1.1 Data set

We obtain flight schedules from the Aviation System Performance Metric (ASPM) for flights landing at LaGuardia International Airport (LGA), New York ,on Feb 17, 2014 [15]. We use this to estimate the arrival demand at LGA on an hourly basis. A discrete time interval of 1 hour is chosen, and all flight times and durations are rounded off to the nearest hour. The arrival demand on Feb 17, 2014 is shown in Figure 4-1(a). The height of the bar is the number of flights that were scheduled to land between the two corresponding x-axis times. For example, between 7 am and 8 am, there were 24 flights scheduled to land. International flights are not considered.

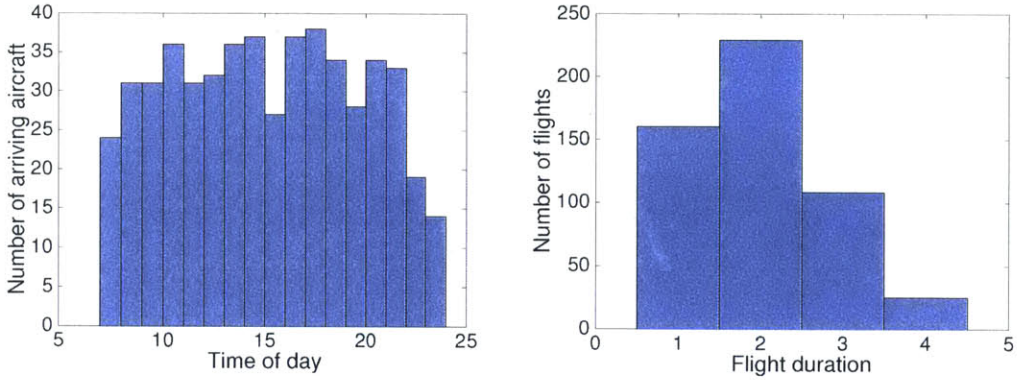

(a) Hourly demand.

(b) Distribution of flight durations.

Figure 4-1: ASPM data for LGA arrivals on Feb 17, 2014.

This is because they are typically not in the scope of a GDP. The distribution of flight times for the 522 flights on that day is presented in Figure 4-1(b).

To summarize, we have the takeoff time, flight duration and landing times for all domestic flights landing at LGA on a particular day. Note that the ASPM database contains information only of carriers that handle atleast 1 % of domestic traffic[15].

4.1.2 GDP parameters

The following parameters characterize the GDPs we simulate

- *The length of the GDP*: In our simulations, we analyze GDPs of 7, 10, and 15 hour durations. In practice, GDPs can be a few hours long, or can even last the entire day. Because so many of them are canceled early or are extended, there is a huge variability in the length of a GDP.

- *The nominal and bad weather capacity*: The nominal capacity of LGA is set to be 10 arrivals/15 min, or equivalently, 40 arrivals/hour. The peak capacity can be as high as 14 arrivals/15 min, but is not sustained for long. We set the reduced capacity to be half the nominal throughput, or 20 arrivals/hour. This is the target arrival rate that is chosen by the GDP planner for the duration of reduced capacity. Operationally, the target arrival rates are lower than what

Prob Index	1	2	3	4	5	6	7	8	9
Mean reduced cap	1.1	1.545	1.96	2.475	3	3.475	3.95	4.425	4.9

Table 4.1: Expected duration of reduced capacity for all the nine probability distributions.

the airport can handle even during reduced capacity. This is done in order to provide buffer for flights that are exempted from GDPs and unscheduled arrivals.

- *The probability of each scenario*: We wish to test the GDP algorithms under a variety of probability distributions. We generate $(2T-1)$ probability distributions for a GDP of length T. Denote $S = \{1, 2, ..., 2T-1\}$ as the set of distributions. Corresponding to every test case $s \in S$, we have an associated probability π_q^s for every branch q in the scenario tree.

- *Ground and air delay cost*: We set $C_a = 2.5$ and $C_{g,1} = 1$.

4.1.3 Probability distributions

We will explain the distributions with the help of an example. Let $T = 5$ and the scenario tree be as shown earlier in Figure 2-2. The nine probability distributions (i.e. $2 \times 5 - 1$) are shown in Figure 4-2.

These $2T-1$ scenarios are such that we get different expected times of capacity improvement. In scenario i, the reduced capacity exists until time i. The expected time of reduced capacity for each scenario is given in Table 4.1. In all subsequent plots in this chapter, we will use the probability scenario index as a surrogate for the expected duration of reduced capacity, for simplicity. A probability scenario of $(2s-1)$ has an expected duration of reduced capacity s.

These distributions can be generalized using the following formula. Define π_q^s as the probability of scenario q happening when we have a probability distribution $s \in S$.

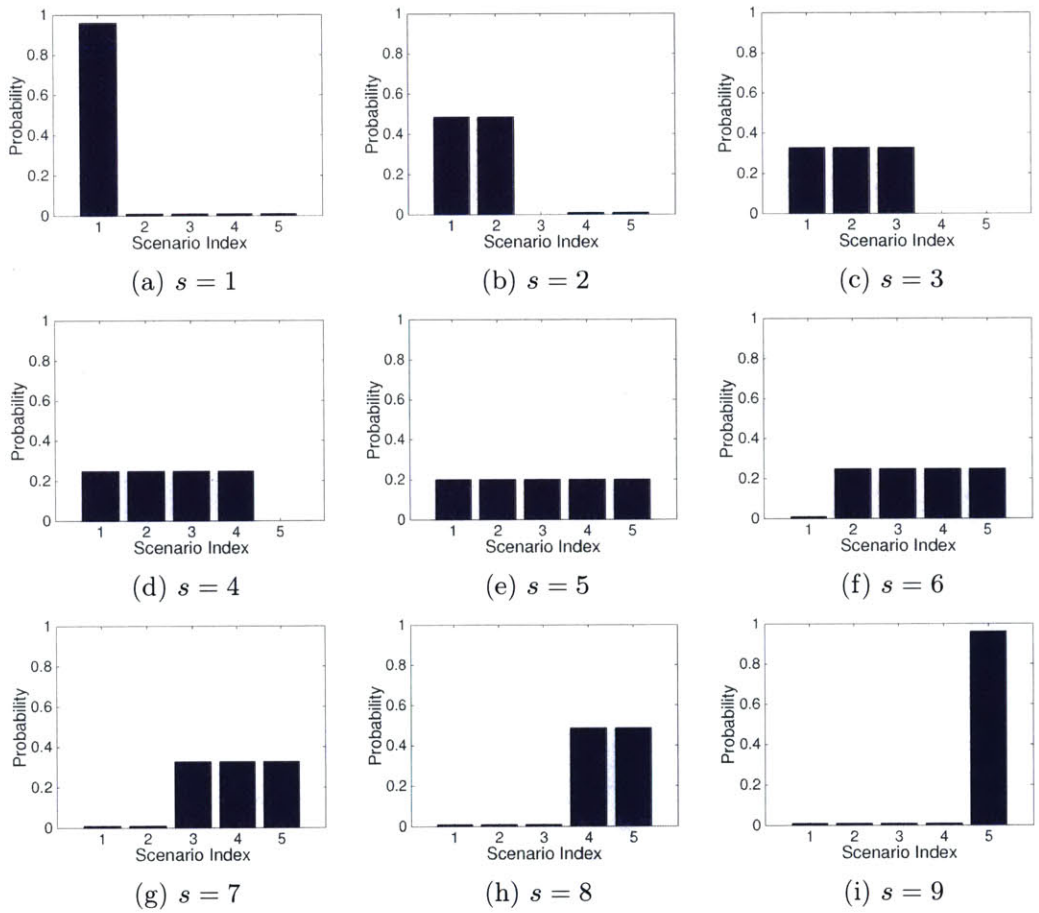

(a) $s = 1$ (b) $s = 2$ (c) $s = 3$

(d) $s = 4$ (e) $s = 5$ (f) $s = 6$

(g) $s = 7$ (h) $s = 8$ (i) $s = 9$

Figure 4-2: The $2T - 1$ probability distributions, with $T = 5$.

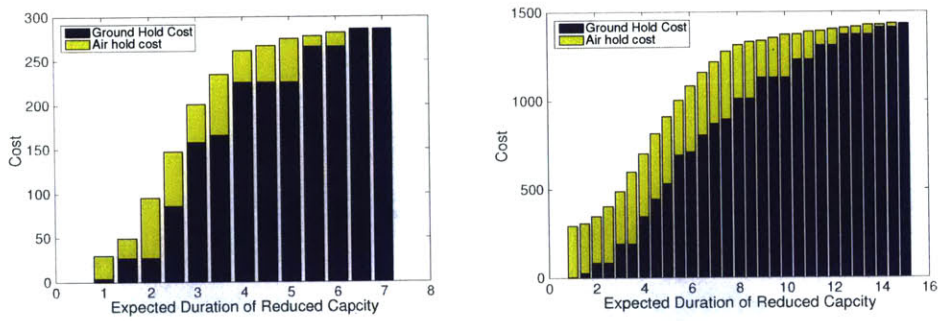

Figure 4-3: Static model costs for $T = 7$ (left), and $T = 15$ (right).

For $s = 1, ..., T$

$$\pi_q^s = \begin{cases} \frac{1-0.01\times(T-s)}{s} & \forall q \in \{1, ..., s\} \\ 0.01 & \forall q \in \{s+1, ..., T\} \end{cases} \tag{4.1}$$

For $s = T+1, ..., 2T-1$

$$\pi_q^s = \begin{cases} 0.01 & \forall q \in \{1, ..., s\} \\ \frac{1-0.01s}{T-s} & \forall q \in \{s+1, ..., T\} \end{cases} \tag{4.2}$$

4.2 Pre-CDM costs for existing models

The static, dynamic and hybrid models are compared in this section. This is the pre-CDM step. We simulate two cases, $T = 7$ and $T = 15$. For each probability distribution, the total cost (which is the sum of the ground hold and the air hold cost) is plotted. The total cost for any given scenario (or equivalently, the expected duration of reduced capacity) follows the order *dynamic* $\leq$ *hybrid* $\leq$ *static*.

In Figures 4-3, 4-4 and 4-5, the pre-CDM cost (ground + air hold cost) is plotted for varying durations of reduced capacity. The air holding cost has a negligible contribution to the total cost in a dynamic model. The dynamic model is able to transfer more air delay to the ground, which makes the total cost low. The static model has the greatest amount of air holding cost, because it is a non-adaptive solution. If

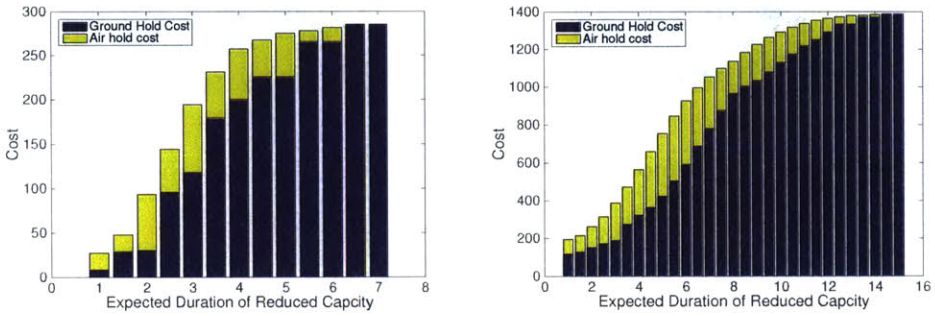

Figure 4-4: Hybrid model costs for $T = 7$ (left), and $T = 15$ (right).

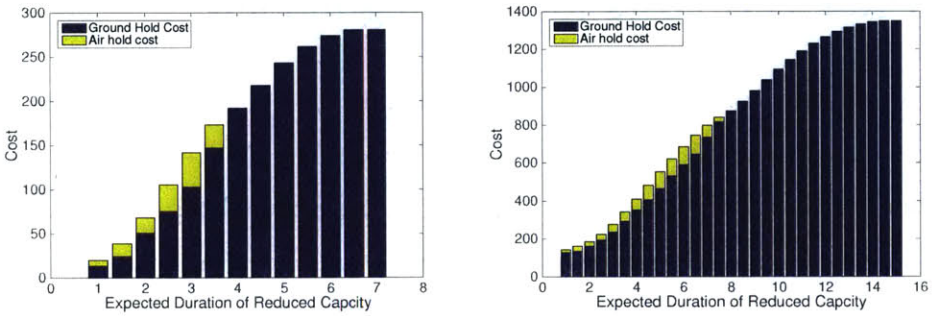

Figure 4-5: Dynamic model costs for $T = 7$, (left) and $T = 15$ (right).

an unlikely scenario with prolonged duration of reduced capacity materializes, then flights will be held in an airborne queue, and incur high costs. The hybrid model has higher cost than the dynamic, but is less expensive than the static.

4.3 RHS model

To implement the RHS model, we need to decide on the number of stages, and the update time of each stage. The greater the number of updates, the closer is the RHS model to a dynamic model. Consequently there is a cost decrease. The goal of this section is to quantify this benefit, and also highlight the need to choose an optimal update time.

First, we present the results for a 2-stage RHS formulation. The update time can be chosen from the set $\{1, ..., T\}$. If the update time is $\tilde{t} = 1$ or $\tilde{t} = T$, the RHS cost will be the same as the static cost. In the former case, no new information is obtained, while in the latter, there are no flights remaining that can make use of the updated information. The plots in Figure 4-6 show the cost as a function of update time. The optimal update time is the one which leads to minimum costs.

This evaluation of optimal update times can be extended to multiple updates as well. For a 3-stage RHS, there will be two update times. When the updates happen at $t = 1$ or $t = T$, or when both happen at the same time ($\tilde{t}_1 = \tilde{t}_2$), there are no benefits of the updates. The cost for two updates $\tilde{t}_1$ and $\tilde{t}_2$ (with $\tilde{t}_2 \geq \tilde{t}_1$) gives us an optimal pair of update times (Figures 4-7 to 4-9).

The optimal update times are different for different probability distributions. Consider the case of T=7, and one update. When the distribution is $s = 1$, the static solution would have come up with an assignment assuming the capacity will improve after $t = 1$. This is because the probability of the capacity improving after $t = 1$ is 0.94. On the other hand, if the capacity does not improve, the static solution will lead to high air hold costs. An update time of $t \geq 2$ is thus useful, and will give valuable capacity information. If the capacity did not clear at the update time, then the second stage can be adjusted accordingly. On the other hand, scenario $s = 7$ will

benefit most from an update which comes later in the GDP, because the probability is spread out and it is hard to predict capacities well in advance. Henceforth, when we use the RHS model, we will always evaluate it using the optimal update times. By this process of enumeration, we can find the optimal update times for any number of stages.

A comparison of the costs for the models across different probability scenarios is summarized in Figure 4-10. As expected, for every probability scenario, having more RHS updates gives lower costs. However, the difference in costs is low, whenever the expected duration of reduced capacity is very low, or very high. They correspond to situations where the probability distributions have the least variance. The information gained through an update is low, and this leads to lower benefits over the static model.

In general, it is observed that there is a decreasing marginal utility to having more updates (Figure 4-11). More than 50% of the cost reduction is obtained with one update itself. Let us consider the case of $s = 1$. The static cost is 29.82, the cost with one update is 21.61, and the lowest achievable cost (with 7 updates) is 20.5. With one update we attain 73% of the benefits. Similarly, for $s = 7$ and $s = 13$, we attain 60.99% and 79.36% of the benefits, respectively. These results justify the use of only one update in the RHS model.

4.4 Intra-airline substitutions

In this section, we simulate the slot substitution process within an airline. Different flight-specific costs motivate airlines to solve an assignment problem where flights are paired with slots, so that the cost is minimized.

For the CDM analysis, we need data on flight specific costs of airlines. Although flight specific costs can be estimated (Ramanujam [25]), these are generally proprietary information and cannot be obtained from public sources. We therefore demonstrate the results qualitatively. We generate a hypothetical test case in which we distribute all the flights into 10 airlines, with equal market share, randomly. The flight-specific air and ground holding costs are derived from a normal distribution.

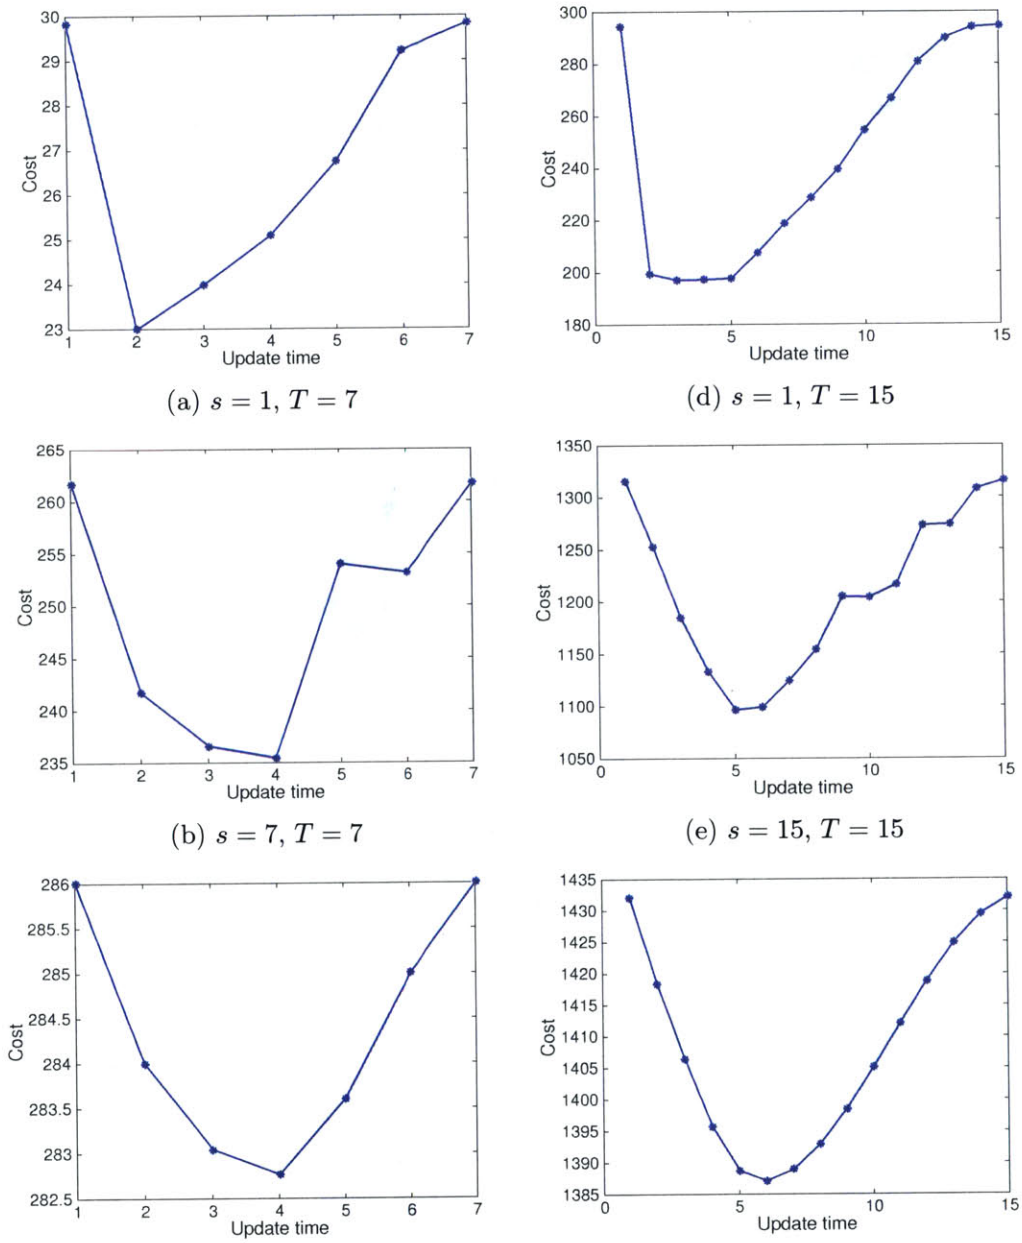

(a) $s = 1$, $T = 7$

(d) $s = 1$, $T = 15$

(b) $s = 7$, $T = 7$

(e) $s = 15$, $T = 15$

(c) $s = 13$, $T = 7$

(f) $s = 29$, $T = 15$

Figure 4-6: One update RHS for $T = 7$ (left), and $T = 15$ (right).

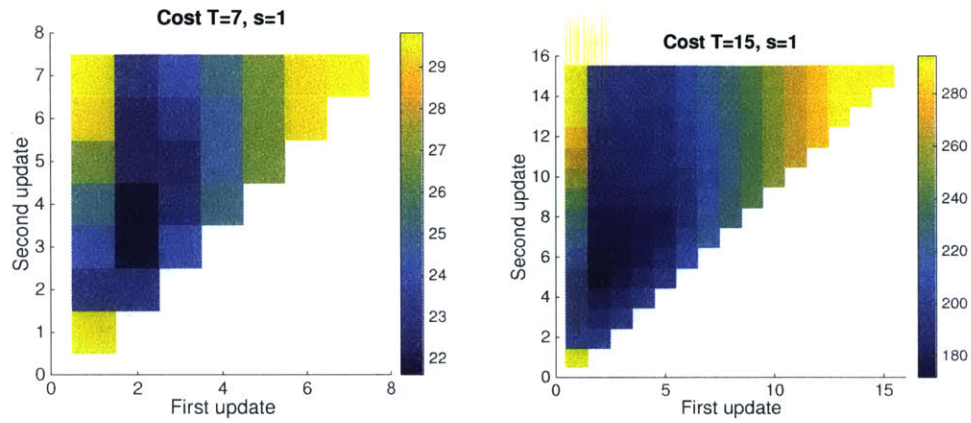

Figure 4-7: Two update RHS costs for $T = 7$ (left), and $T = 15$ (right) with $s = 1$.

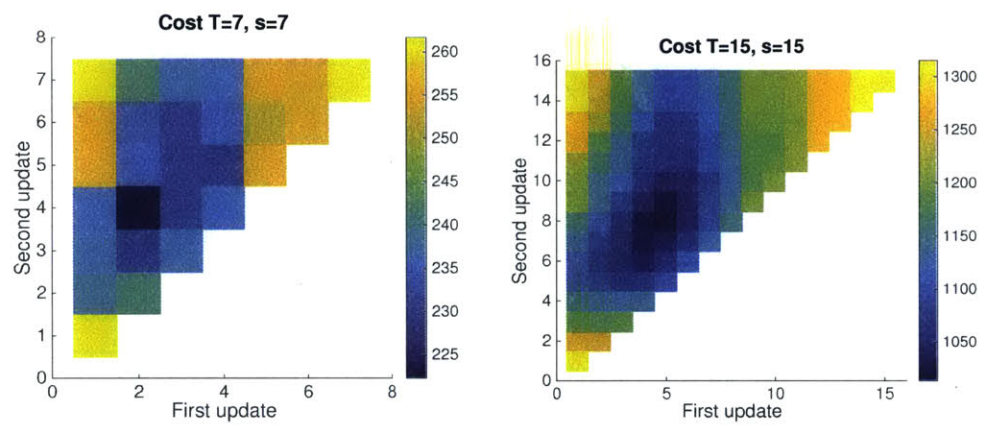

Figure 4-8: Two update RHS costs for $T = 7$ (left), and $T = 15$ (right) with $s = T$.

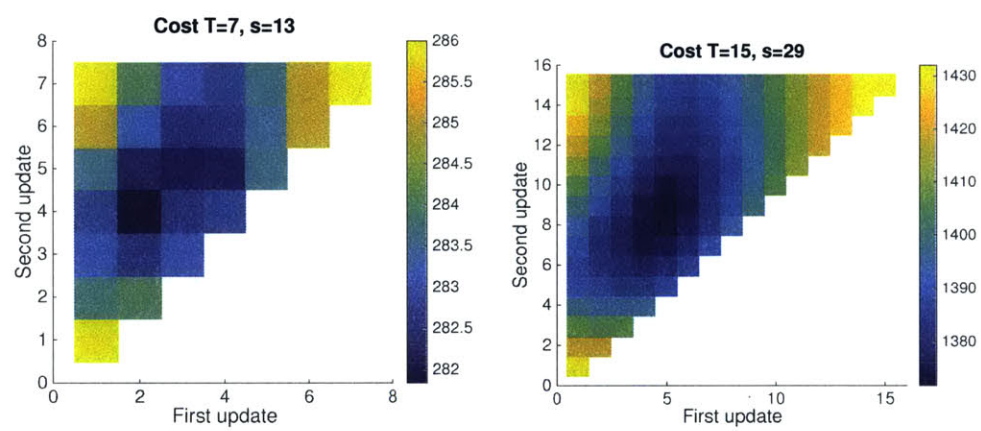

Figure 4-9: Two update RHS costs for $T = 7$ (left), and $T = 15$ (right) with $s = 2T-1$.

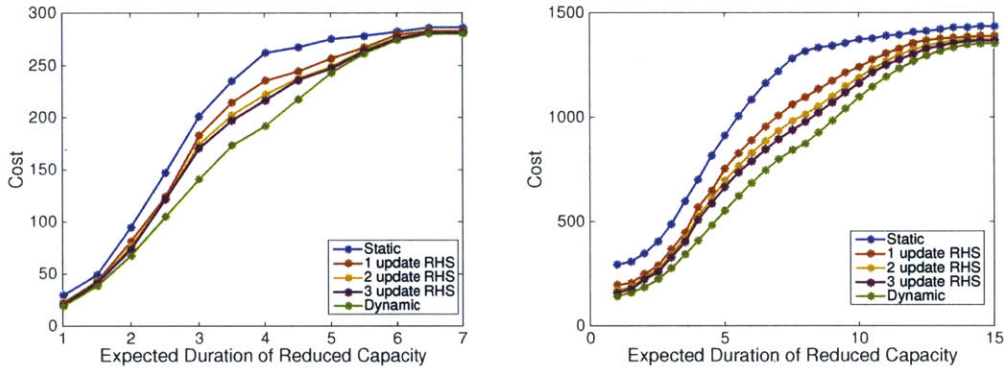

Figure 4-10: Cost of the RHS model in comparison to the static and dynamic model, for $T = 7$ (left), and $T = 15$ (right).

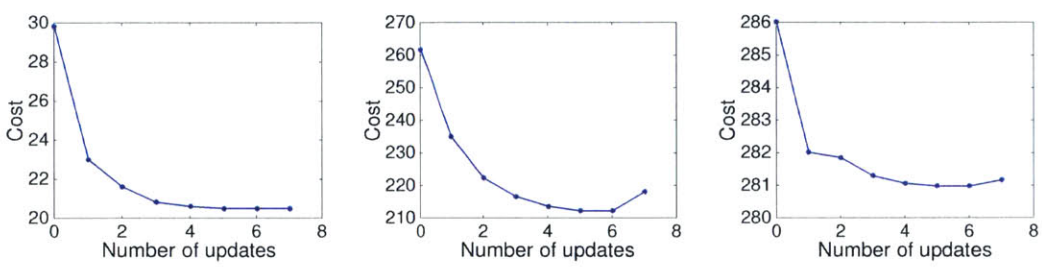

Figure 4-11: Optimal cost for different number of updates. Zero updates means that it is a static model. The different probability scenarios plotted are $s = 1$ (left), $s = 7$ (center), and $s = 13$ (right). $T = 7$ for all the three cases.

The mean ground holding cost is assumed to be 1 and the mean air holding cost is assumed to be 2.5. The coefficient of variation is 25% for both the air and ground hold costs.

For a GDP with T=7, the results after the swap are shown in Figure 4-12. When the expected duration of reduced capacity is low, the swaps do not change the rankings of the models. The static model has the highest cost and the dynamic model has the lowest. As the expected delay increases, intra-airline swaps start playing a greater effect. The number of flights getting rescheduled increases. There is greater opportunity to decrease costs due to swaps, and this benefits the static and RHS models more than the dynamic model. The dynamic model ultimately has the highest costs, and the static has the least. The performance of the RHS model is interesting: It is never the worst performing model. In Figure 4-12 (right), we show the percentage improvement of every model over the worst. The missing bars represent the worst model. The RHS model outperforms the worst model by 25% in some cases.

When T=15, the trends are similar (Figure 4-13). Due to the longer duration of the GDP, the differences in performance of the four models gets amplified. The hybrid model performs much better and is comparable to the RHS model, which is again the most consistent model. The difference between choosing an optimal and a non-optimal model can be as high as 50%. The choice of a GHP model is important, and we demonstrate that the RHS is consistent across a variety of probability distributions and GDP lengths. In the next section, we also show that the RHS model is robust to different levels of cost variations.

4.5 Effect of cost variation

It is to be expected that an airline with more variation in flight-specific costs will benefit more from the swap process. This is clear when comparing Figure 4-14 and Figure 4-15. The Coefficient of Variation (CV) in costs is a measure of dispersion, or spread, in the flight-specific costs across all the flights for a given airline. If the CV is higher, there is more variability in the flight-specific cost for an airline. If the CV

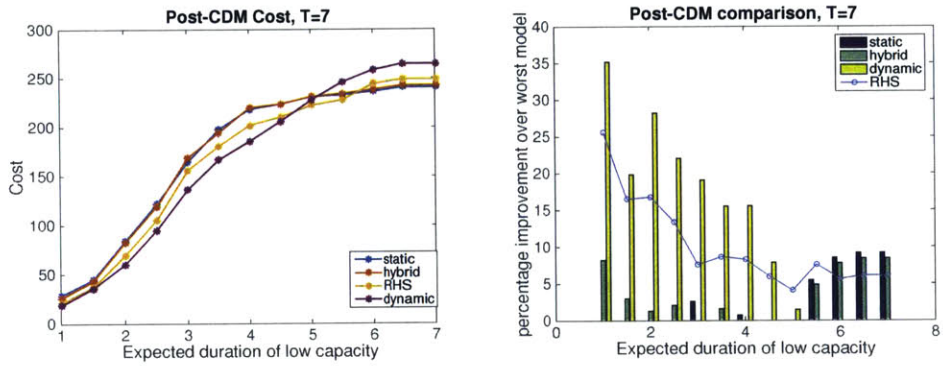

Figure 4-12: Post-CDM cost of all the four models for $T = 7$ (left), and the percentage improvement over the worst model (right). The missing bars correspond to the model with the highest cost.

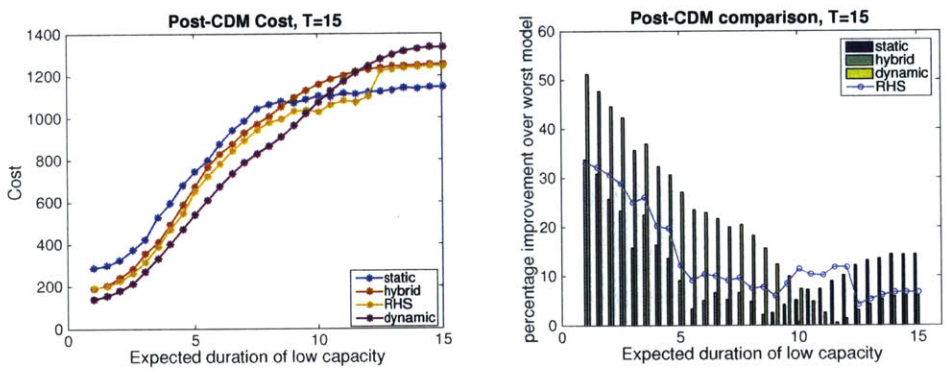

Figure 4-13: Post-CDM cost of all the four models for $T = 15$ (left), and the percentage improvement over the worst model (right). The missing bars correspond to the model with the highest cost.

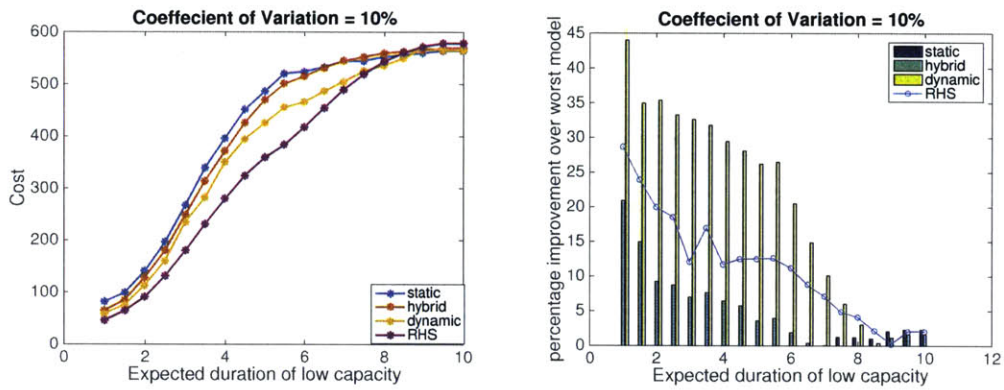

Figure 4-14: Actual costs (left), and percentage improvement over the worst model (right), for $T = 10$, coeff. of variation $= 10\%$.

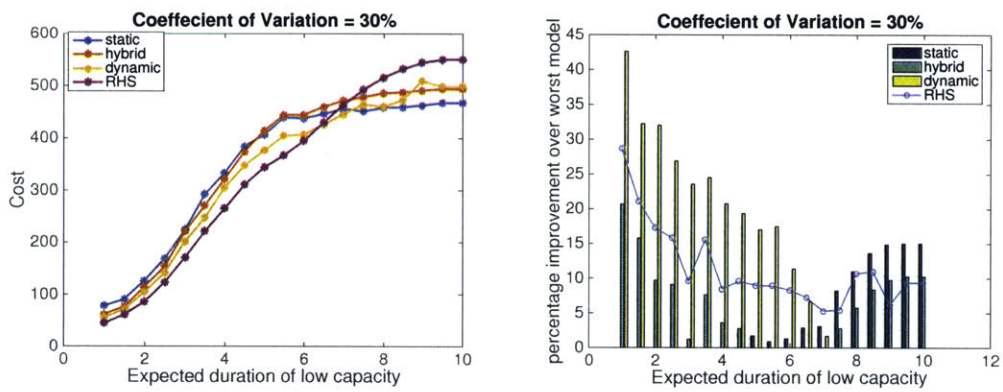

Figure 4-15: Actual costs (left), and percentage improvement over the worst model (right), for T=10, coeff. of variation $= 30\%$.

is lower, all the flights will have a similar cost for getting delayed. As the coefficient of variation in costs increases, the static, hybrid and RHS models benefit greatly. Naturally, these differences are more prominent when the GDPs are longer, and a large number of flights get swapped.

4.6 Slot credit substitution

An airline is also allowed forfeit a slot ($slot_c$), and request an alternative slot at time k. We denote quantities pertaining to the canceled flight using a subscript c. Although

a slot is requested at time k, it is not always possible to accommodate the request. The objective of the SCS optimization is to minimize this deviation from k, while respecting the feasibility constraints. There are three indicators of the performance of the SCS mechanism:

1. *Cost of canceled flight*: When a flight cannot be granted a landing slot at time k, there is an additional ground holding cost. There could also be an airborne queue preventing the flight from landing at k. The term B_c is the additional cost that a canceled flight has to incur. $B_c = G_c + A_c$, where G_c is the ground holding cost and A_c is the air holding cost.

2. *Benefit to other flights*: When a flight is rescheduled, others are compressed, and have a decrease in delay costs. The sum of this benefit for other non-canceled flights is B_o.

3. *System benefits*: The system-wide benefit from the SCS request is the difference between B_o and B_c, i.e, $B_s = B_o - B_c$.

When different GHP models are used, there is a varying amount of flexibility to assigning a new slot for an aircraft. This causes variations in the cost B_c. Intuitively, the static model can assign any requested slot k without any additional ground holds, but there might be airborne delays, which the static model cannot avoid. On the other hand, the dynamic model will have reduced flexibility in allotting the requested slot k (since they are duration specific), but manages air hold costs better. In terms of B_o, a similar argument holds, and the static assignment can be compressed more easily in order to transfer the benefits of the cancellation to other flights. We will quantify these benefits through simulations.

We use a test data set, instead of operational data, to simulate SCS requests. We consider a GDP with T=10, and an arrival demand of $(16, 16, 15, 16, 15, 16, 16, 15, 8, 5)$. These 138 flights are randomly assigned a flight duration between 2 to 5 hours. The nominal and reduced capacity are 16 arrivals/hour and 8 arrivals/hour, respectively. The costs are averaged over 100 SCS triggers for each probability distribution. Flights

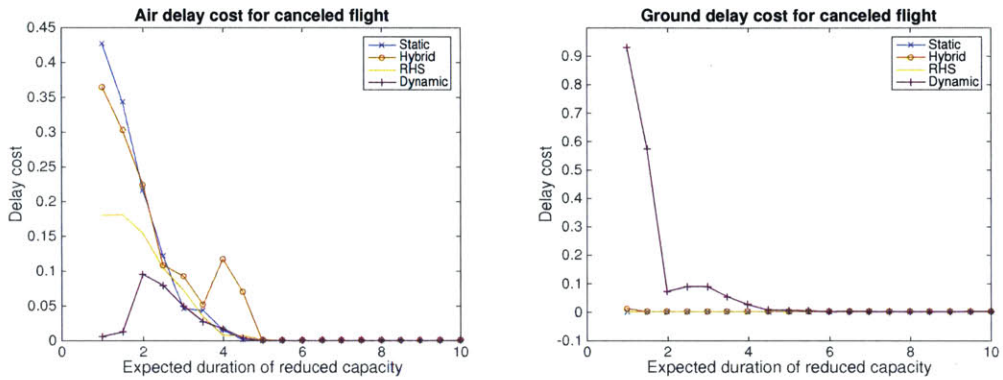

Figure 4-16: Additional air holding (left), and ground holding (right) costs for the canceled flight.

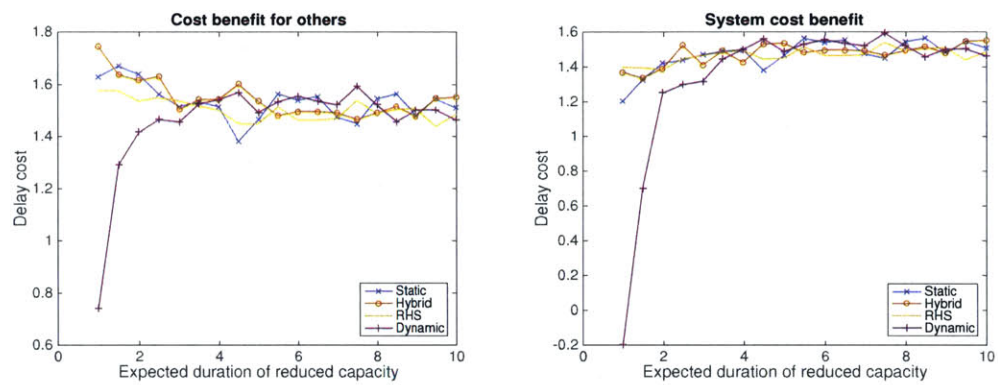

Figure 4-17: Mean benefit for other flights (left), and for the system (right), when an SCS is triggered.

arriving between $t = 1$ and $t = 5$ are canceled randomly, and they request a slot that is 1 to 4 hours beyond the scheduled time. The ground and air holding costs are held the same for all flights, with $C_{g,1} = 1$ and $C_a = 5$.

The plots of A_c and G_c is shown in Figure 4-16. Note that the scales of the y-axis on both plots are different. Despite the differences in allocation flexibility, all four models are similar over a wide range of capacity disruptions. When G_c is zero, it means that the requested slot k is allotted to the aircraft. All models except the dynamic are able to achieve this. The dynamic model, because of limited swap flexibility, is unable to always allot the compensatory slot at time k. However, the

air holding cost is the least for the dynamic model. Finally, when the GDPs have a long duration of reduced capacity, there is no difference between the four models.

The flexibility to interchange slots determines the benefit to other flights (B_o). The simulations (Figure 4-17) indicate that there is no significant difference among the models in this regard.

The dynamic model, in a scenario where the capacity improves quickly (expected duration of reduced capacity close to 1), can lead to negative system benefits, B_s. The high B_c ends up decreasing the system benefits. For all other cases, no model outperforms the other. The RHS model is therefore comparable to the existing models, in terms of the SCS performance.

4.7 Discussion

The RHS model, based on a simple idea of repeatedly running the static model, performs well under a variety of scenarios. The pre-CDM cost of the RHS model is more than that of the dynamic, but less than that of the static model. Furthermore, it is observed that even a single update is sufficient to capture most of the benefits of a multistage static model. This is important from an operational point of view, and also to maintain slot swap flexibility. In the pre-CDM step, the 2-stage RHS shows a comparable performance to the hybrid model, while eliminating the need for scenario specific slots. Consequently, the RHS is the most consistent performer in the intra-airline swap stage. It also shows favorable results in the SCS procedure.

To summarize, the RHS model is found to be a viable alternative to the hybrid model, as it incorporates CDM procedures while designing the optimization framework. In terms of implementation, it involves less information overhead, results in more predictable operations, and provides an equitable distribution of delays among airlines.

Chapter 5

Conclusions

With air traffic projected to grow significantly, congestion in airports is an increasing concern. A long-term solution to the problem of congestion is to increase the capacity at airports. This involves investments in airport infrastructure, and will take a long time to implement. At a more tactical level, significant improvements in system efficiency can be made through the better management of existing resources.

The Ground Delay Program (GDP) is a commonly used strategy to manage periods of significant congestion. Developing efficient and fair ground holding strategies will reduce costs, decrease fuel burn and improve safety. In this thesis, we draw ideas from two existing GHP solution strategies to develop a more robust alternative. The new receding horizon model is operationally easy to implement, more robust and considers the collaborative decision making step in its design.

The 2-step RHS model is tested using operational data, and is found to be a viable alternative to existing GDP models. As the RHS model is built on multiple runs of the static model, it also has favorable equity properties. The allocation follows the Ration By Schedule (RBS) principle where the order of the flights in the original schedule and the revised allotment are same. The RHS model thereby confirms to acceptable notions of fairness among the airlines.

There are two issues that must be kept in mind while implementing the proposed framework. The first is related to the scenario tree. Weather reports need to be made available in this particular probabilistic format for running the model. The second

assumption is regarding information exchange. The revised schedule is time-sensitive information that should be made available to all stakeholders at the update time. The revision might require airlines to make immediate decisions at the update time, such as ensuring a particular flight departs in the next 1 hour. Maintaining appropriate ground support, crew availability, and adequate time for passenger boardings is critical, and is not accounted for in our model. A possible alternative, if it is not possible to share revised schedules at the update time, is to allow airlines to follow the full results of the initial static solution.

There are several promising extensions to this work:

- Solving the GHP for finer time-intervals (15 min, for example) is important from an operational perspective. Preliminary simulations indicate that the results are similar to the 1-hour interval results. The computational complexity of the static, hybrid and RHS models do not grow exponentially with $|T|$. However, the dynamic model scales in a nonlinear fashion with the planning duration T, leading to greater computational expense.

- A more relaxed notion of fairness could be considered for the SCS mechanism. Currently we require every flight to experience lower air and ground holding times after compression. Airlines might like to have benefits at an aggregate level that includes all their flights, rather than individual flights. This relaxation will lead to greater flexibility in slot compression, pushing the final solution closer to a global optimum from the current Pareto-optimal solution.

- The benefits of having multiple rounds of intra-airline slot substitutions or scenario specific slot substitutions may also be explored. In practice, airlines can re-evaluate the delay sensitivity of flights, be impacted by another Traffic Management Initiative elsewhere in the system, or learn new information (for example, regarding crew availability) only later during the GDP. The potential benefits of making another round of substitutions under such scenarios are not obvious. It is also not known how such a concept of multiple static updates will extend to the multi-airport ground holding problem.

- The models described in this thesis have two steps. The first step is a central allocation by the FAA while the second step involves decentralized optimization by airlines. This two-step solution paradigm tries to incorporate two key features of the current system: First, the airlines have access to private cost information which they do not share. Second, it is important to be fair to different airlines and flights of different durations. Relaxing these two constraints has the potential to lower costs. Let us suppose that the airlines report the delay cost of each of their individual flights truthfully to the central planner. It is then possible to set up a single-step optimization that can incorporate flight-specific costs. Such an optimization may not necessarily result in a fair allocation across airlines. For example, larger aircraft with more passengers are likely to have larger delay cost and will receive priority over smaller aircraft. A single-step optimization can be formulated for the static, dynamic, RHS and hybrid formulations and it is expected that the single-step dynamic model will be the best performer. This extension will enable the computation of a globally optimal solution to the GDP. The ratio of the two-step cost to the globally optimal cost is a measure of the price of fairness in this system. It reflects the additional delay costs, or the price that the entire system pays for being fair. Studying single-step GHP formulations can help develop better ground holding policies. A related question pertains to the incentive compatibility of the mechanism. Would airlines benefit by reporting truthfully? Would certain airlines misreport their delay costs in order to avoid delays? Analyzing the cost-reporting process from a mechanism-design perspective will propel these one-step models closer to practical implementation.

- The effect of unplanned aircraft arrivals during GDPs is an aspect that has received little attention in previous work [11]. Dynamic models can be reframed so that the current queue length is also incorporated at each update time in order to model random aircraft arrivals. It is worth investigating how the performance of the other three models, namely, the static, hybrid and RHS,

would deteriorate for different levels of unplanned arrivals.

Bibliography

[1] Hamsa Balakrishnan. Techniques for reallocating airport resources during adverse weather. In *Decision and Control, 2007 46th IEEE Conference on*, pages 2949–2956. IEEE, 2007.

[2] Michael O Ball, Robert Hoffman, and Avijit Mukherjee. Ground delay program planning under uncertainty based on the ration-by-distance principle. *Transportation Science*, 44(1):1–14, 2010.

[3] Michael O Ball, Robert Hoffman, Amedeo R Odoni, and Ryan Rifkin. A stochastic integer program with dual network structure and its application to the ground-holding problem. *Operations Research*, 51(1):167–171, 2003.

[4] Michael O Ball, Robert L Hoffman, Amedeo Odoni, and Ryan Rifkin. The static stochastic ground holding problem with aggregate demands. 1999.

[5] Jonathan F Bard and Dinesh Natarajan Mohan. Reallocating arrival slots during a ground delay program. *Transportation Research Part B: Methodological*, 42(2):113–134, 2008.

[6] Cynthia Barnhart, Dimitris Bertsimas, Constantine Caramanis, and Douglas Fearing. Equitable and efficient coordination in traffic flow management. *Transportation Science*, 46(2):262–280, 2012.

[7] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250, 2012.

[8] Dimitris Bertsimas and Shubham Gupta. Fairness in air traffic flow management. In *INFORMS Meeting, San Diego, CA, USA*, volume 32, page 7, 2009.

[9] Dimitris Bertsimas and Amedeo Odoni. A critical survey of optimization models for tactical and strategic aspects of air traffic flow management. 1997.

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

[11] Narender Bhogadi. Modeling demand uncertainties during ground delay programs. Master's thesis, 2002.

[12] Lorenzo Brunetta, Guglielmo Guastalla, and Lisa Navazio. Solving the multi-airport ground holding problem. *Annals of Operations Research*, 81:271–288, 1998.

[13] Jonathan Cox and Mykel J. Kochenderfer. Optimization approaches to the single airport ground hold problem. In *Journal of Guidance, Control, and Dynamics*, Kissimmee, FL, 2015.

[14] Jonathan Cox and Mykel J. Kochenderfer. Ground delay program planning using markov decision processes. *Journal of Aerospace Information Science*, 2016.

[15] Federal Aviation Administration. ASPM database, 2014. `https://aspm.faa.gov/apm/sys/main.asp`.

[16] Eugene P Gilbo. Airport capacity: Representation, estimation, optimization. *Control Systems Technology, IEEE Transactions on*, 1(3):144–154, 1993.

[17] Charles N Glover and Michael O Ball. Stochastic optimization models for ground delay program planning with equity–efficiency tradeoffs. *Transportation Research Part C: Emerging Technologies*, 33:196–202, 2013.

[18] Robert Hoffman and Michael O Ball. A comparison of formulations for the single-airport ground-holding problem with banking constraints. *Operations Research*, 48(4):578–590, 2000.

[19] Balázs Kotnyek and Octavio Richetta. Equitable models for the stochastic ground-holding problem under collaborative decision making. *Transportation Science*, 40(2):133–146, 2006.

[20] P Liu and Mark Hansen. Scenario-free sequential decision model for the single airport ground holding problem. In *Proceedings of the 7th USA/Europe Air Traffic Management R&D seminar*, 2007.

[21] Bengi Manley and Lance Sherry. Analysis of performance and equity in ground delay programs. *Transportation Research Part C: Emerging Technologies*, 18(6):910–920, 2010.

[22] Avijit Mukherjee and Mark Hansen. A dynamic stochastic model for the single airport ground holding problem. *Transportation Science*, 41(4):444–456, 2007.

[23] Avijit Mukherjee, Mark Hansen, and Shon Grabbe. Ground delay program planning under uncertainty in airport capacity. *Transportation Planning and Technology*, 35(6):611–628, 2012.

[24] Gordon F Newell. Airport capacity and delays. *Transportation Science*, 13(3):201–241, 1979.

[25] Varun Ramanujam. *Estimation and tactical allocation of airport capacity in the presence of uncertainty.* PhD thesis, Massachusetts Institute of Technology, 2011.

[26] Varun Ramanujam and Hamsa Balakrishnan. Estimation of arrival-departure capacity tradeoffs in multi-airport systems. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 2534–2540. IEEE, 2009.

[27] Varun Ramanujam and Hamsa Balakrishnan. Increasing efficiency though flexibility in airport resource allocation. In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pages 1759–1766. IEEE, 2014.

[28] Octavio Richetta and Amedeo R Odoni. Solving optimally the static ground-holding policy problem in air traffic control. *Transportation Science*, 27(3):228–238, 1993.

[29] Octavio Richetta and Amedeo R Odoni. Dynamic solution to the ground-holding problem in air traffic control. *Transportation Research Part A: Policy and practice*, 28(3):167–185, 1994.

[30] Mostafa Terrab. *Ground holding strategies for air traffic control.* PhD thesis, Massachusetts Institute of Technology, 1990.

[31] Thomas Vossen and Michael Ball. Optimization and mediated bartering models for ground delay programs. *Naval Research Logistics (NRL)*, 53(1):75–90, 2006.

[32] Thomas Vossen, Michael Ball, Robert Hoffman, and Michael Wambsganss. A general approach to equity in traffic flow management and its application to mitigating exemption bias in ground delay programs. *Air Traffic Control Quarterly*, 11(4):277–292, 2003.

[33] Thomas WM Vossen and Michael O Ball. Slot trading opportunities in collaborative ground delay programs. *Transportation Science*, 40(1):29–43, 2006.

[34] Peter B Vranas, Dimitris J Bertsimas, and Amedeo R Odoni. The multi-airport ground-holding problem in air traffic control. *Operations Research*, 42(2):249–261, 1994.

[35] Peter BM Vranas, Dimitris Bertsimas, and Amedeo R Odoni. Dynamic ground-holding policies for a network of airports. *Transportation Science*, 28(4):275–291, 1994.