# An Analytic Study of the Benefits of Collaborative Arrival Planning (CAP) on Air Carrier Ground Operations

by

Kari Andersson

and

William Hall

Charles Stark Draper Laboratory

May 24, 2000

# Contents

# List of Figures

# List of Tables

# Chapter 1

# Introduction

The purpose of this study is to explore the operational benefits of increased communication and collaboration between the airlines and controllers during the arrival process. Through Collaborative Arrival Planning (CAP), airlines would have access to regularly updated arrival time estimates for their aircraft from the Traffic Management Advisor (TMA) and the Final Approach Spacing Tool (FAST). Additionally, with some enhancements, CAP could allow airlines to communicate priority levels of their arriving aircraft. This information could then be incorporated into the decisions made by the controllers when sequencing the arrival stream.

Via CAP, CTAS is already providing American Airlines and Delta Airlines at the Dallas Fort Worth Airport (DFW) with more accurate landing time estimates from TMA in real-time [12]. CAP could be enhanced to accept arrival sequence preferences from the airlines and transfer that information to CTAS. By altering the algorithms in TMA, CTAS could provide to the air traffic controllers sequencing advisories that consider the airline arrival sequence preferences. This approach would have little impact on the controller's current workload; indeed, it should reduce workload somewhat.

In order to test the benefits of offering these functions with CAP, we built the Aircraft Sequencing Model (ASM), an optimization model that simulates airline operational decisions about aircraft movement times under resource constraints. The model considers arrival sequence, departure schedule, physical gate resource and ground crew resource constraints in determining an arrival sequence that minimizes passenger delay. The potential benefits are quantified by testing several scenarios using historical data.

Many data sets are input into ASM for this analysis. A detailed description of the use of the data is included in Chapter 3. Historical data are used to determine scheduled and actual wheels-on (ON), gate arrival (IN) and gate departure (OUT) times. These times are used to define feasible movement times as well as to calculate delays. Airline-specific data are used to determine ground resource availability and capacity as well as to determine ground resource needs by aircraft. To prioritize aircraft and to estimate passenger connection data, the number of passengers on each aircraft and the number of passengers connecting between aircraft are simulated. Finally, three parameters were tuned for the model: the taxi-in time for each aircraft, the delay costs and the sequencing flexibility the airlines are allowed.

Using these data, ASM generates sequence preferences to be used in an arrival sequencing tool, such as CTAS. It also produces estimated aircraft ground movement times which can be used by the airline in its ground operations decisions. Figure 1-1 provides a schematic of the data flows of ASM.

Results indicate that the airline's operational efficiency would be improved using these CAP-provided tools. More accurate landing time estimates could reduce passenger delays up to 3%, saving up to 2000 passenger minutes of delay in a 3.25 hour time period. Depending on the level of flexibility allowed in reordering the aircraft, preferential sequencing could reduce passenger delays 5–20%, saving over between 700,000 and 1.8 million passenger minutes of delay per month.

This document is intended to provide a detailed description of the ASM optimization model and a deeper exploration of the inputs into and the outputs from the model. Chapter 2 provides an overview of the design of the model, including critical assumptions and a description of how the model is used to generate results. Chapter 3 provides a description of all data sources used and a detailed explanation of the manipulations done to the data. Chapter 4 details the formulation of the model, including a discussion of the boundary effects caused by the finite time horizon considered and a description of the effort to reduce the number of constraints and variables considered. Chapter 5 includes a discussion of the

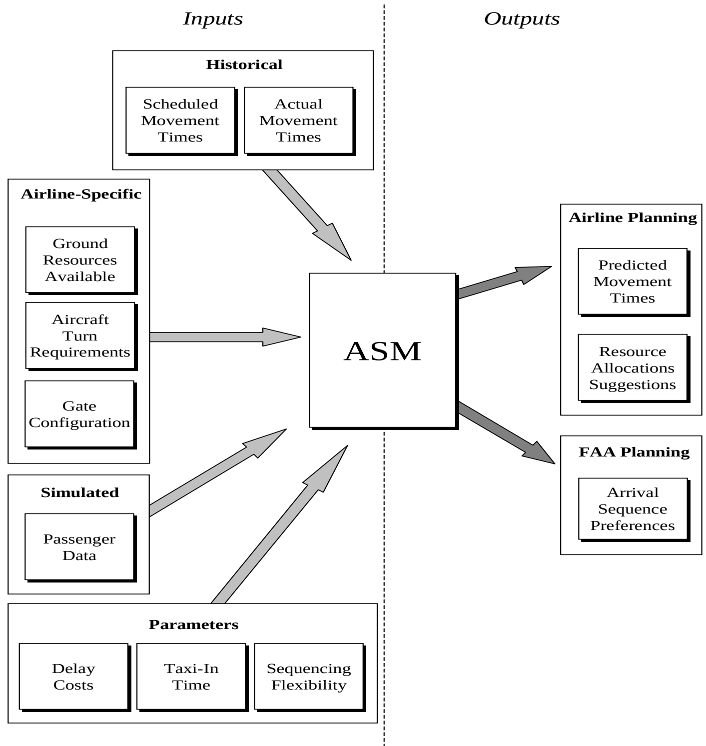


Figure 1-1: Information flows in ASM

calibration and validation process for some of the major assumptions in the model. Further, it contains a validation of the model's results. A discussion of the results can be found in Chapter 6. Chapter 7 concludes the document with a summary of the results.

# Chapter 2

# Modeling Approach

The objective of this study is to quantify the benefits of using CAP to share information between airlines and air traffic controllers during the arrival process. The information to be shared includes expected landing times of arriving aircraft and airline arrival sequence preferences. This study tests for potential benefits such as more efficient use of ground resources, reduced delays affecting passengers, and reduced departure delays affecting the entire system. These benefits are quantified from the airline perspective, but the benefits extend to the entire air transportation system.

In order to capture the benefits to the airline, we model decisions affected by this exchange of information, namely all ground operations and aircraft movement decisions. In particular, we model the aircraft's landing, gate arrival and push back times, accounting for taxi-in time, gate assignment constraints and ground resource constraints. The movement times are then determined to minimize delays, including gate availability delay (time an aircraft waits for a gate after taxiing), arrival delay (minutes from scheduled arrival), departure delay (minutes from scheduled departure) and missed connection delay (minutes until a passenger can make a connecting flight, when the original connection is missed). In order to quantify the benefits to the airline, scenarios are run on the model to compare the delays incurred under varying assumptions.

This chapter provides an overview of the Arrival Sequencing Model (ASM) developed to quantify these potential benefits. The first section provides a description of the model.

The second section describes the scenario analyses used to generate results. Finally, the simplifying assumptions incorporated in the design of ASM are discussed.

## 2.1 Model Description

Many factors are incorporated into the design of ASM, including gate constraints, ground resource constraints, airline fairness, and the types of airports and time periods appropriate for the model.

Gate space is a finite resource at any airport and requires efficient management. Aircraft are assigned to gates hours in advance of their arrival. However, like the rest of the air transportation system, delays and mechanical failures can disrupt this schedule. The assignment of aircraft to gates is further complicated because each gate is designed to service only certain types and sizes of aircraft. Additionally, sometimes large aircraft can render an adjacent gate unusable. Gate availability affects arrival and hence, push back times.

There are numerous operations involved in turning an aircraft, including de-planing and boarding passengers, unloading and loading baggage, catering, cleaning, refueling, and maintenance. Each of these operations is handled by a distinct set of ground resources. However, not all of these resources are critical to the turn time of an aircraft[1]. Key airline operations personnel indicated during interviews that baggage unloading and loading is often a bottleneck in the turning process. Therefore, the only ground crew explicitly considered in the initial design of ASM are baggage handlers.

In an industry with thin profit margins [1], airlines compete for market share to drive up their profitability. Consequently, airlines are sensitive to policy or procedural changes that benefit one airline more than another. ASM considers airline fairness in its design, meaning ASM guarantees that an airline does not improve its operational performance at the expense of another airline's. Fixing the airline's landing times in the model enforces airline fairness; an airline is allowed to shuffle aircraft landing times only within its set of input landing

[1] Turn time is defined as the time between arrival at the gate and push back from the gate. A short turn time leads to higher utilization of resources, which leads to improved return on assets for the airline.

times.

The exchange of information between airlines and air traffic controllers will be most beneficial at hub airports. At hub airports, there is usually a bank structure to the schedule, meaning a wave of arrivals land and a corresponding wave of departures take off about an hour later, allowing passengers to connect to the departing flights. During these waves, there are many aircraft to sequence and many passengers to connect. Further, hub airports tend to be more ground resource constrained due to the bursts of activity during arrival and departure waves that require many people, equipment and gates. Therefore, ASM is best suited for time periods covering these banks of arrivals and departures. This study examines one hub airport for a particular airline. This airport will be referred to as Airport 1 from here on.

ASM makes decisions about aircraft movement times respecting the above considerations. It explicitly makes trade-offs between delaying aircraft and maximizing passenger connectivity. Therefore, the model considers operational delays: gate availability delay, arrival delay and departure delay. Further, it considers passenger delays, or the delay caused by a missed connection. Therefore, the objective function of the model is designed to minimize operational and passenger delays, as measured in passenger minutes.

## 2.2 Results Approach

Recall that the objective of the study is two-fold: to estimate the benefits of more accurate arrival time predictions and to estimate the benefits of sequencing flexibility in the arrival process. We use ASM to achieve these objectives.

The value of more accurate arrival time predictions is estimated using a perturbation analysis. Assuming that there is no flexibility in arrival time (i.e., the airline is not allowed to change the sequence of its arriving aircraft), the set of feasible landing times is perturbed slightly and the effect on the system measured. The significance of these effects is tested using standard statistical methods.

The value of preferential sequencing is measured by assuming a fixed set of feasible landing times, but altering the level of sequencing flexibility. Sequencing flexibility is represented by the minutes an aircraft is able to move up or back its landing time (within the set of input landing times). The significance of these effects is tested using standard statistical methods. This analysis does not account for limits on the level of sequencing flexibility in the system. Instead, its aim is to measure the potential benefits, leaving the testing of the feasibility of the resulting sequences to another study.

## 2.3 Simplifying Assumptions

No model can fully capture all system effects. This section outlines the major simplifying assumptions underlying ASM. Later sections discuss the impact of some of these assumptions on the model's solution.

First, the model assumes that all arrival sequences, within the set of input landing times, are feasible. However, this is not necessarily true. For example, if A and B are two aircraft flying sequentially through the same TRACON fix, it is unlikely that the order of these aircraft can change before landing. Since this study is intended to identify whether providing opportunities for preferential sequencing is beneficial, this simplification is acceptable. Further studies will be necessary to incorporate such constraints on arrival sequences.

Second, the model is designed to consider aggregated physical gate constraints. To manage the size of the model, in terms of the number of variables and constraints, ASM does not assign each aircraft to a particular gate. Instead, the model assigns the aircraft to *some* gate, assuming that gate supply constraints aggregated by aircraft type are met. It is possible that the model produces a solution for which there is no feasible gate assignment given the arrival at gate and push back times.

Further, ground resources are modeled such that the number of resources assigned can vary over time. For example, an aircraft could be assigned three baggage handlers every other minute. Additionally, ASM allows the baggage handlers to be assigned to *any* aircraft. In practice, baggage handlers are assigned to zones, or a subset of gates, and a small team of handlers is assigned to an aircraft while it is parked at one of the gates in the zone. As a result, ASM could produce a solution for which no team can be assigned continuously.

The limitations of ASM regarding the ground operations constraints (gates and baggage handler assignments) have not shown to impact the validity of the model. Feasibility tests were performed on a set of ASM solutions, as described in Section 5.3.

ASM is a determininstic model, meaning there is no stochasticity incorporated in its design. In particular, the taxi-in time of each aircraft is assumed constant. Analyses in [2] indicate that the taxi-in process at airports is highly stochastic. This implies that ASM does not capture some important system dynamics on the ground. However, interviews with key airline personnel indicate that a constant taxi-in time is currently assumed by the airlines when making decisions about ground operations. Therefore, the ASM's constant taxi-in time assumption is no worse than current operating assumptions for the airlines. The process leading to this taxi-in assumption is described in Section 5.1.

Finally, the objective function is measured in passenger-minutes of delay, which is not a metric directly linked to the airline's cost structure. The translation of this metric to dollars is difficult. For example, consider a passenger rebooked on a flight departing four hours after the scheduled departure time of the original connection. These 240 minutes of passenger delay are not easily translated into costs (or lost revenues) to the airline. Sometimes the delay will have no impact; there may be no direct cost to the airline from the delay and the passenger may not rebook her flight on another airline. On the other hand, sometimes the delay will have an impact; the airline may have to provide room and board for the delayed passenger (cost) or the airline may have to provide the passenger a ticket on a competitor's flight (cost and lost revenue). However, passenger minutes of delay does link both operational efficiency and the passenger experience, both of which have an effect on the profitability of the airline. Therefore, it provides an effective metric for the purpose of this study.

# Chapter 3

# Inputs and Data Issues

Since ASM considers ground operations at an airport, there is a significant amount of data that is needed to make the model reflect true operations. The main data source for this analysis is the Airline Service Quality Performance (ASQP) data set, which records push-back from gate (OUT), take-off (OFF), landing (ON), and arrival at gate (IN) times. From this data, we estimate actual landing, minimum turn and taxi-in times, as well as scheduled arrival and departure times. Additional data sources are used by the model. Airport-specific data are necessary for scheduled turn times, resources scheduled to manage the turn operations and physical gate constraints. Equipment data for each aircraft are necessary to link the ASQP data and the airport-specific data. To capture downstream effects of delay, a metric is used to translate current delay into downstream delay. And finally, data on the number of connecting passengers is necessary for accurately assessing the importance of aircraft connections.

This section provides descriptions of the data sources used in the study. Furthermore, it covers how the data were manipulated in order to incorporate them into ASM.

## 3.1 Historical Movement Data

The analyses discussed herein heavily rely on the Airline Service Quality Performance (ASQP) database, which provides information about the jet operations of 10 major airlines: Alaska, American, America West, Continental, Delta, Northwest, Southwest, TWA, United and US Airways. For most of these airlines' flights, ASQP provides actual push-back (OUT), take-off (OFF), landing (ON) and gate arrival (IN) times, as recorded by Aircraft Communication Addressing and Reporting Sensors (ACARS) and processed by Aeronautical Radio, Incorporated (ARINC) [10]. If one of the above-mentioned ten airlines have jet aircraft not equipped with ACARS, it is required to provide these data elements manually.

This aircraft movement data is then linked to the Official Airline Guide (OAG), which contains planned flight times for all scheduled air carrier and commuter flights, and the Computer Reservation System (CRS), which contains scheduled flight times, possibly updated from the previously released OAG schedule. The full ASQP database consists of the following data: aircraft tail number, airline, flight number, origin, destination, date, arrival time (actual, OAG scheduled, CRS scheduled), pushback time (actual, OAG scheduled, CRS scheduled), actual landing time, actual take-off time, taxi-in and taxi-out times, arrival and departure delay, actual and scheduled ground-to-ground time and airborne time.

The ASQP database includes only data relating to jet aircraft. This means that not all flights at an airport are captured in the database. Further, some flight legs are missing from the data, meaning that an aircraft, identified by tail number, may arrive or depart without a corresponding departure or arrival. Comparisons between the actual operations counted at DFW and those captured in the ASQP database reveal that the ASQP database is relatively complete with respect to jet aircraft. In July 1998, the ASQP database included on average about 469 departures for American and about 128 for Delta per day. DFW reports an average of 470 and 128 jet departures for American and Delta, respectively [5]. The ASQP database is essentially complete in terms of jet aircraft operations. However, DFW further reports about 248 and 74 propeller aircraft departures for American and Delta, respectively [5]. Therefore, the ASQP database excludes about 35% and 37% of these airlines' departure traffic at DFW. The extent to which data is excluded depends on the usage of propeller planes at an airport. However, it is reasonable to assume that other hub airports will have similar completeness levels.

The accuracy of the ASQP data was confirmed with independent observations. Visual observations at Boston Logan airport confirmed ASQP recorded push-back times [4]. Take-off and landing times recorded in the ASQP data closely matched estimated take-off and landing times generated from high-resolution, timed radar tracks provided by CTAS at the DFW airport [2].

### 3.1.1 Feasible Landing Times

In order to address the issue of airline fairness, a set of landing times is input into ASM that defines the landing times the model can assign to an aircraft. ASM finds an optimal one-to-one mapping of landing times to aircraft. In effect, the model finds the optimal assignment of aircraft to landing times. We refer to this input set of landing times as the feasible set of landing times for the airline. For the analyses conducted in this study, the set of feasible landing times is the set of historical landing times for the airline. This assumption means that if the model assigns an aircraft an earlier landing time (relative to its actual), it can only do so by assigning another aircraft a later landing time (relative to its actual).

Physical limitations of the aircraft, such as acceleration and speed, and safety considerations, such as separation, restrict the set of possible sequence changes. Rather than explicitly incorporate these limitations into ASM, we fix the magnitude of the adjustment with an input parameter $Window$. In other words, ASM is designed such that it can assign an aircraft a landing within $Window$ minutes of its input landing time, which for the analyses in this study is the actual landing time. We refer to this set of possible landing times for an aircraft as the set of feasible landing times for the aircraft.

For example, assume the set of feasible landing times for the airline is given by

$$FeasLand = \{\ldots, 10{:}00,\ 10{:}07,\ 10{:}08,\ 10{:}10,\ 10{:}11,\ 10{:}15, \ldots\}$$

and $Window = 5$. Consider aircraft $ac$ with the actual landing time of 10:08. The set of feasible landing times for aircraft $ac$ is given by

$$FeasLand_{ac} = \{10{:}07,\ 10{:}08,\ 10{:}10, 10{:}11\}.$$

The impact of the parameter $Window$ on the results is addressed in Chapter 6.

### 3.1.2 Scheduled Departure Times

In theory, an aircraft should not push back before its scheduled departure time. In practice, however, this is not the case. At Airport 1 in March 1998, 39% of the departing planes left before scheduled departure time. To accommodate for this in the model, the scheduled departure times for the aircraft that pushed back early, are set equal to the actual departure times. However, an aircraft is not considered late unless it pushes back after the originally scheduled departure time. In effect, instead of having one "cost-free" departure time, each aircraft that departed early is assigned a continuous set of cost-free departure times ranging from its actual departure time to its original scheduled departure time. The departure delay cost function is described in detail in Section 3.4.

### 3.1.3 Taxi-In Times

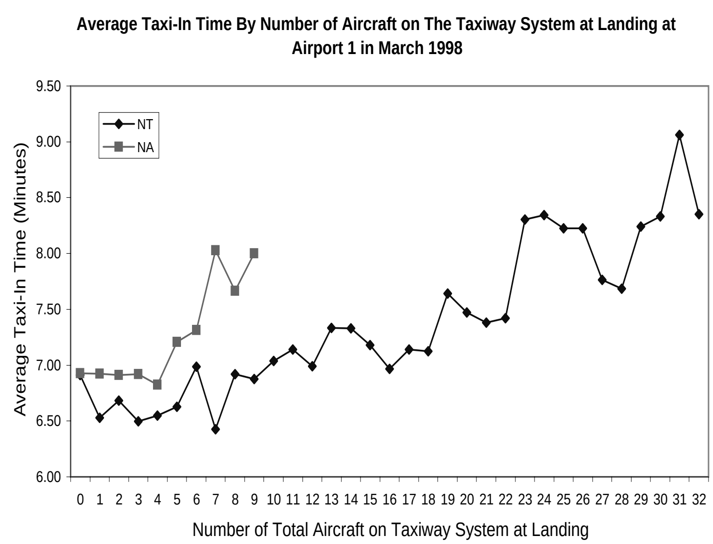


Figure 3-1: Average taxi-in time as a function of congestion

Taxi-in time is defined as the time from landing (ON) to arrival at the gate (IN). The ON to IN time consists of two components: the taxi time to the gate from the point of landing and any delay at runway/taxiway intersections or at the gate. From the ASQP data it is impossible to distinguish the two components. As suggested in [10], an estimate of the true taxi-in time can be made by considering the "unimpeded taxi-in time", defined as the taxi-in time when there are few active aircraft on the ground to impede the arriving aircraft's movement. This results in an unimpeded taxi-in time distribution.

If we let $NA_t$ and $NT_t$ be the number of arriving aircraft and total aircraft, respectively, on the taxiway system at the time of landing, $t$, then we can plot the taxi-in time distribution as a function of these two congestion measures. For Airport 1 we see in Figure 3-1 that the mean taxi-in time increases as both $NA_t$ and $NT_t$ increase, with $NA_t$ resulting in a sharper increase. This increase is caused, in part, by the increased taxi time to the gate under congested conditions, and in part by the increased queuing at runway/taxiway crossing and at the gate. Of these potential causes, only gate congestion is considered in the model.

ASM assumes the taxi-in times for each aircraft is given. This assumption is made because predicting taxi-in times as well as aircraft movement times requires an iterative solution method. As shown in Figure 3-1, there is a strong correlation between taxi-in time and congestion level. This correlation implies that a tool that predicts taxi-in time should consider congestion level. In order to determine congestion levels, the movement times of the aircraft must be known. Therefore, in order for ASM to accurately predict taxi-in times, the aircraft movement times would be assumed fixed. Then given the taxi-in time predictions, the aircraft movement time predictions would be updated. This iterative procedure would continue until the solutions converged. This approach is impractical for the purposes of this study. Therefore, the taxi-in times are assumed known *a priori*.

As discussed in Section 5.1, a number of taxi-in time assumptions were tested during the calibration process. The tests indicate that the truest-to-actual results occur when ASM assumes a constant taxi-in time set to the overall average taxi-in time. The extent to which this assumption affects the model is not included in this study. However, interviews with key airline personnel indicate that the airlines currently assume a constant taxi-in time when making decisions about ground operations. Therefore, the taxi-in time assumption in ASM is no worse than the current operating assumption.

## 3.2 Airport-Specific Data

Many of the assumptions in the model are specific to the particular airline and airport examined. These include physical gate constraints, scheduled turn time, ground resource staffing levels by aircraft type, and the total number of available ground resources. The data incorporated into the model are based on those used in the airline's daily operations, obtained from interviews with key airline personnel at the airport.

### 3.2.1 Gate Configuration Data

The gate constraint data consist of physical restrictions at each gate. This includes the types of aircraft that can be serviced at each gate, as well as any adjacent gate restrictions. For example, gate 12A can service all types of aircraft, but when servicing a widebody, adjacent gate 12 becomes unusable.

Recall that physical gate constraints are aggregated by aircraft type in ASM. Further, adjacent gate restrictions are not included in the model. Section 5.3 explains these aggregation assumptions in more detail and provides some insight into how they affect the solution.

### 3.2.2 Turn Data

The turn data consist of the minimum time and ground resources scheduled for turning an aircraft. Interviews with key airline personnel suggested that the baggage handling process was a bottleneck of the turn process. Therefore, ASM models baggage handlers. The number of baggage handlers scheduled to work a plane depends on the type of aircraft and the origin-destination pair for the aircraft. For example, a widebody from an international origin turning to a domestic destination will have a different minimum turn time and ground resource allotment than a narrow-body from a domestic origin turning to a domestic destination.

Multiplying the minimum scheduled turn time by resource allotment data determines a required number of resource-minutes for turning the aircraft. For example, if the minimum scheduled turn time for an aircraft is 40 minutes and the number of baggage handlers allotted to that aircraft is 3, then a total of 120 resource-minutes is required to turn the aircraft. This required number of resource-minutes is what is used in ASM. The transformation is necessary to allow the model to assign more resources to an aircraft to reduce its turn time.

To increase the accuracy of the model, future versions of ASM might consider explicitly modeling other ground crew resources, such as mechanics, gate agents, caterers, etc. Examining the trade-offs between problem size and solution accuracy will help determine which resources to include in the model.

There are certain situations where an airline may decide to remove an aircraft from a gate. For example, if an aircraft needs mechanical work, the airline may decide to do the repairs off the gate in order to keep the gate available for other incoming aircraft. It is impossible to determine such events precisely from the ASQP database. However, these events usually imply a relatively long turn time. Therefore, an aircraft is considered to have left the gate (for something other than a departure) if the turn time is greater than 3 hours.

If an aircraft is removed from the gates, an estimate of the time at which it returns to the gates is needed. Since the minimum scheduled turn time for the aircraft is known, the aircraft is assumed to be brought to the gate the minimum scheduled turn time *plus* a buffer of 5 minutes prior to the scheduled departure time. So, if an aircraft is scheduled to depart at 1:20 and its minimum scheduled turn time is 40 minutes, ASM assumes it is at the gates at 12:35.

If the time the aircraft is assumed to be at the gate is before the beginning of the time horizon considered in the model, an adjustment on the required resources is made. This adjustment is based on the ratio of the length of time the aircraft has been at the gate prior to the first time period in the model relative to the length of time scheduled to be at the gate. In the above example, if the time horizon begins at 1:00, the required resources are reduced by a factor of $25/45 = 0.55$, meaning if the required resources were 120 resource-minutes, the adjusted required resources are $120 \times (1 - .55) = 54$ resource-minutes.[1]

[1] As discussed in Section 4.5, some aircraft included in the model are at the gate at the beginning of the time horizon. The adjustment on the required resources for these aircraft is the same as for aircraft with turn times greater than 3 hours, but the actual arrival at gate time is used instead of the estimated arrival at gate time. If the aircraft has been at the gate longer than its minimum scheduled turn time (indicated

There are further limitations to the assignment of resources in ASM. First, the number of resources assigned to an aircraft at any point in time is bounded from above by the scheduled allotment of baggage handlers plus two. This bound was confirmed in interviews with key airline personnel. Second, the minimum turn time of an aircraft is bounded from below by the actual minimum turn time. The actual minimum turn time takes into account that the process of de-planing and boarding passengers has a positive lower bound. The actual minimum turn time for each aircraft type was estimated using historical turn times recorded in the ASQP database.

## 3.3 Equipment data

Much of the airport-specific data is based on equipment type of the aircraft, while the ASQP data identifies an aircraft by its tail number. In order to link these two sets of data, the Aviation Gold Mine's Aircraft Owners database, compiled by Software Innovations, which contains equipment information for US registered aircraft, was used. The database contains the FAA record of registration for an aircraft.

This data set is relatively complete. Table 3.1 shows the percentage of aircraft that are included in the August 1998 ASQP database but *not* included in the Aircraft Owners database for each airline at the George Bush International Airport in Houston (IAH) and the Chicago O'Hare Aiport (ORD). Notice that for American Airlines (AA), the percentage of missing aircraft is quite high. However, this database is more complete than any other equipment source tested. We currently have no explanation for the fact that so many of American Airline's aircraft can not be found in existing databases.

In the event that the equipment type is missing for an aircraft, its equipment type is assigned randomly based on the distribution of aircraft types that are flown to the same destination. This is similar to the approach taken in [7].

by a ratio $> 1$), the number of required resources is set to 0.

| Airline | IAH | | | ORD | | |
|---|---|---|---|---|---|---|
| | Total Aircraft | Missing Aircraft | Percent Missing | Total Aircraft | Missing Aircraft | Percent Missing |
| CO | 332 | 10 | 3.01 | 216 | 6 | 2.78 |
| AA | 260 | 47 | 18.08 | 588 | 177 | 30.10 |
| NW | - | - | - | 249 | 5 | 2.01 |
| UA | 177 | 0 | 0.00 | 483 | 0 | 0.00 |
| US | 130 | 0 | 0.00 | 229 | 2 | 0.87 |
| DL | 211 | 0 | 0.00 | 332 | 0 | 0.00 |

Table 3.1: Completeness of equipment database

## 3.4 Departure Delay Costs

Departure delays can affect more than the passengers onboard the aircraft; there are potential downstream effects on the airline's operations. The aircraft and the crew aboard might need to make connections at the destination airport. If the aircraft is delayed, these connections might be missed, disrupting the schedule. These disruptions can be very costly to the airline.

In order to estimate the magnitude of these downstream effects American Airlines and Oak Ridge National Laboratory in [3] determined a Delay Multiplier (DM). The authors claim that the full delay to the operations in the system due to an initial delay of an aircraft can be estimated by multiplying the initial delay by the DM. Since the analysis done to determine the DMs considered only crew and aircraft sequences, disregarding the impact on passengers, cargo and gate space, the DMs are considered by the authors of [3] to be conservative estimates. The DMs are determined by the magnitude of the initial delay and the time of day the delay occurs. Long delays early in the day have a large multiplier effect. Similarly, early delays have a greater multiplier effect than later delays.

The DM is incorporated in the cost of departure delays considered in ASM. The magnitude of the delay is determined by the model and the timing of the delay is based on the scheduled departure time. Recall that the objective function is measured in passenger-minutes. The departure delay cost considered in ASM, therefore, is the product of the departure delay, the DM, and the number of departing passengers. This assumes that the number

| Aircraft Size | Direct Operating Costs ($) | Direct Operating Costs (mins) |
|---|---|---|
| Medium | 2.5 | 5 |
| Large | 4.5 | 9 |
| Heavy | 6.0 | 12 |

Table 3.2: Direct operating costs for aircraft waiting to depart

of passengers on the aircraft affected downstream is similar to the number of passengers on the initially delayed flight. Considering that most crews can operate only specific aircraft types, it seems reasonable that the flight to which the crew is connecting has an aircraft of a similar size. Assuming, then, that the load factors are similar for downstream flights, crew and aircraft sequence delays would affect a similar number of passengers, justifying our cost coefficient.

The DOT's measurement of airline on-time performance is measured as the percentage of departures that push-back from the gates within 15 minutes of the scheduled departure time. In other words, an airline is not "penalized" for holding a departure up to 14 minutes after its scheduled departure time. To capture this effect in the model, a different cost function is used for the first few minutes after scheduled departure time. The cost function during these few minutes represents the direct operating costs of holding an aircraft on the ground. The objective function coefficient calibration process, as discussed in Section 5.1, indicated that the best-fit results occur when the direct operating costs are considered for the first 9 minutes after scheduled departure time.

According to [11], an airline incurs costs while an aircraft ready for departure sits on the ground. These costs, representing the cost of the crew aboard the aircraft, are included in Table 3.2. Recall that the objective function is measured in passenger minutes. To translate these direct operating costs from dollars to minutes of delay, 60 minutes of passenger delay is approximated to be equivalent to $30. This yields the direct operating costs per minute in units of minutes of delay shown in Table 3.2.

As discussed in Section 3.1.2, some aircraft actually depart *before* scheduled departure time. For these aircraft, the cost function for the time period between actual departure time and scheduled departure time is negative. In particular, the cost associated with departing $d$ minutes early is $-d$. In other words, the function is such that aircraft are nominally encouraged to depart early. However, the negative cost is such that the benefit of leaving early will never be greater than the benefit of holding an aircraft for passenger connections.

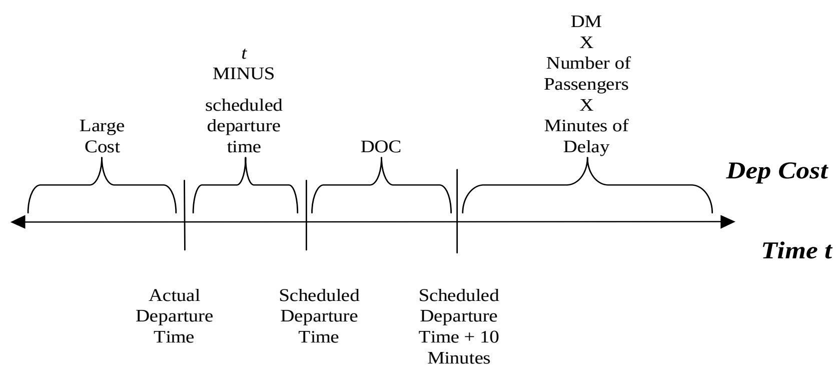


Figure 3-2: Departure cost function over time

Figure 3-2 illustrates the break-down of the departure cost function. The calibration and validation of the cost function is described in Section 5.2.

## 3.5 Passenger Data

Another important parameter of the model is the number of passengers on each flight and the number of passengers connecting between flights. Both of these are estimated probabilistically according to the approach designed by Hall in [7]. The approach is detailed in [7], but a brief description of the process is provided here. In this approach, the number of passengers on a particular aircraft is based on a normal distribution centered at 75% of the number of seats on the aircraft, with a standard deviation of 25% of the number of seats on the aircraft. Once the number of passengers on each aircraft is determined, the number of connecting passengers on the aircraft is determined using a binomial distribution, where parameter $pcf$ denotes the probability that a passenger makes a connection. The flight to which the passenger is connected is probabilistically chosen from a set of feasible connections. A validation of this approach is included in [7].

# Chapter 4

# Formulation

Using the approach and data described above, this section details the formulation of ASM assuming a finite time horizon of between three and four hours. The first four sections define the variables, parameters, constraints and objective function, respectively. The fifth section introduces some boundary conditions and the impact these conditions have on the formulation. Finally, methods for managing the number of variables and constraints included in ASM are discussed.

## 4.1 Variables

There are five main sets of variables used in the model. The indices over which these variables are defined are based on the set *Planes*, which is the set of all aircraft included in the model, and the set *Time*, which is the set of all time units considered in the model. The variables are described in Table 4.1 below.

An important part of the success of this formulation is the definition of the three movement variables as having value one at time $t$ if the aircraft moved *by* time $t$ instead of *at* time $t$. This variable definition led to frequent integer optimal solutions to the relaxed problem and in the cases where the optimal solution to the relaxed problem was not integral, very few nodes in the branch and bound tree were examined before an optimal integer solution was found.

| Variable | Definition |
|---|---|
| $Land_{ac,t}$ | binary, value 1 if aircraft $ac \in Planes$ lands *by* time $t \in Time$ |
| $Gate_{ac,t}$ | binary, value 1 if aircraft $ac \in Planes$ arrives at a gate *by* time $t \in Time$ |
| $Dep_{ac,t}$ | binary, value 1 if aircraft $ac \in Planes$ pushes back *by* time $t \in Time$ |
| $Connect_{ac1,ac2}$ | binary, value 1 if passengers from aircraft $ac1 \in Planes$ can connect to aircraft $ac2 \in Planes$ |
| $NumRes_{ac,t}$ | indicates the number of ground resources assigned to aircraft $ac \in Planes$ at time $t \in Time$ |
| $Delay_{ac}$ | indicates the number of minutes of gate availability delay for $ac \in Planes$ |

Table 4.1: Description of variables

## 4.2 Parameters

There are a number of key parameters in this formulation. The determination of most of the parameters is discussed in Chapter 3. Table 4.2 includes a list of the parameters and their definitions.

## 4.3 Constraints

The model consists of numerous sets of constraints. The constraints are given here, followed by more detailed explanations.

| Parameter | Definition |
|---|---|
| $Window$ | number of minutes an aircraft can move up or back its landing time |
| $NumLand_t$ | total number of aircraft that can land at time $t$ |
| $Taxi_{ac}$ | minutes of "unimpeded" taxi-in time assigned to aircraft $ac \in Planes$ |
| $Turn_{ac}$ | minimum turn time for aircraft $ac$ |
| $Connect$ | number of minutes it takes a passenger to connect to a departing aircraft |
| $ResNeeds_{ac}$ | minimum number of resource-minutes required to turn aircraft $ac \in Planes$ |
| $Res$ | total number of resources available during any minute of time period |
| $ResMax_{ac}$ | maximum number of resources that can be assigned to aircraft $ac \in Planes$ |
| $SizeAC_{s,ac}$ | indicator variable which is 1 if aircraft $ac \in Planes$ is an aircraft of size $s$ |
| $NumGates_s$ | total number of gates available for aircraft of size $s$ |
| $TotGates$ | total number of gates available |

Table 4.2: Description of parameters

$$
\begin{aligned}
Land_{ac,t} - Land_{ac,t+1} &\leq 0 \quad \forall\, ac,\ t && (4.1)\\
Gate_{ac,t} - Gate_{ac,t+1} &\leq 0 \quad \forall\, ac,\ t && (4.2)\\
Dep_{ac,t} - Dep_{ac,t+1} &\leq 0 \quad \forall\, ac,\ t && (4.3)\\
Land_{ac,T} &= 1 \quad \forall\, ac && (4.4)\\
Gate_{ac,T} &= 1 \quad \forall\, ac && (4.5)\\
Dep_{ac,T} &= 1 \quad \forall\, ac && (4.6)\\
\sum_{ac} (Land_{ac,t} - Land_{ac,t-1}) &\leq NumLand_t \quad \forall\, t && (4.7)\\
Land_{ac,t-Taxi_{ac}} - Gate_{ac,t} &\geq 0 \quad \forall\, ac,\ t && (4.8)\\
Gate_{ac,t-Turn_{ac}} - Dep_{ac,t} &\geq 0 \quad \forall\, ac,\ t && (4.9)\\
Delay_{ac} - \sum_{t=1}^{T} (Land_{ac,t} - Gate_{ac,t}) &= -Taxi_{ac} \quad \forall\, ac && (4.10)\\
Cnx_{ac1,ac2} - Gate_{ac1,t} + Dep_{ac2,t+Connect} &\leq 1 \quad \forall\, ac1,\ ac2,\ t && (4.11)\\
\sum_{i=1}^{t-1} NumRes_{ac,i} & && (4.12)\\
-Dep_{ac,t} \times ResNeeds_{ac} &\geq 0 \quad \forall\, ac,\ t && \\
NumRes_{ac,t} - ResMax_{ac} \times Gate_{ac,t} &\leq 0 \quad \forall\, ac,\ t && (4.13)\\
\sum_{ac} NumRes_{ac,t} &\leq Res \quad \forall\, t && (4.14)\\
\sum_{ac} (Gate_{ac,t} - Dep_{ac,t}) \times SizeAC_{s,ac} &\leq NumGates_s \quad \forall\, s,\ t && (4.15)\\
\sum_{ac} (Gate_{ac,t} - Dep_{ac,t}) &\leq TotGates \quad \forall\, t && (4.16)
\end{aligned}
$$

The first three sets of constraints arise from the definition of the three movement variables. Recall that these are binary variables with value 1 if the aircraft $ac$ has moved *by* time $t$. Constraints (4.1)–(4.3) enforce the value of the variable to increase monotonically over time, namely, once the value of the variable is 1, it remains 1 for the remaining time periods in the model.

The second three sets of constraints, (4.4)–(4.6), are cover constraints. Each aircraft must land, arrive at the gate, and push-back in the time horizon. This equates to forcing the value of the variables to 1 in the last time period, $T$.

To ensure airline fairness, the model allows the airline to swap aircraft landing times only within its set of feasible landing times. This set is represented by the parameter $NumLand_t$, which indicates the number of landings possible at time $t$. With (4.7), the model restricts the number of landings at time $t$ to be the same as the input landing times.

There is a natural sequence to events at an airport. First, the aircraft lands. Then it taxis and arrives at the gate. After it is processed, or "turned", it is pushed back from the gate. To enforce this sequence of events, the model includes Constraints (4.8) and (4.9). Constraint (4.8) enforces that an aircraft can not arrive at the gate until it has landed and taxied to the gate. Similarly, Constraint (4.9) ensures that an aircraft can not push back from the gate until it has arrived at the gate and been turned. Note that an aircraft $ac$ is considered on the taxiway system at time $t$ if it has landed but not yet arrived at the gates, namely if $Land_{ac,t} - Gate_{ac,t} = 1$. Similarly, an aircraft $ac$ is considered at the gates at time $t$ if it has arrived at the gates but not yet pushed back from the gates, namely if $Gate_{ac,t} - Dep_{ac,t} = 1$.

Recall that gate availability delay is the minutes that an aircraft sits on the airport surface waiting for a gate after taxiing. Gate delay is defined as the arrival time minus the landing time minus the taxi-in time of the aircraft, as indicated in Constraint (4.10).

Constraint (4.11) is equivalent to the constraint used by Hall in [7]. This constraint tests whether passengers on arriving aircraft $ac1$ could connect to departing aircraft $ac2$. Recall that $Connect$ is the minimum number of minutes required for an arriving passenger to connect to a departing flight. Therefore, this constraint says for potentially connecting aircraft $ac1$ and $ac2$, if $ac2$ departs *before* $ac1$ arrives or $ac2$ departs less than $Connect$ minutes *after* $ac1$ arrives, then passengers from $ac1$ cannot connect to $ac2$.

Constraints (4.12)–(4.14) are concerned with the resources assigned during the turning of an aircraft. Recall that the ground resources considered in the model are baggage handlers. Constraint (4.12) restricts the departure time of an aircraft to be after all the baggage handler-minutes required for turning the aircraft have been assigned. Constraint (4.13) restricts the number of baggage handlers assigned to any particular aircraft at any point in time to be less than the maximum number of baggage handlers that can be assigned to the aircraft. Furthermore, it ensures resources are not assigned before an aircraft has arrived at the gate. Finally, Constraint (4.14) restricts the number of baggage handlers assigned to all aircraft at any point in time to be less than the total number of baggage handlers available.

The last two sets of constraints represent the physical gate configuration. As discussed in Section 3.2.1, gate constraints are modeled in aggregate, meaning an aircraft is not assigned to a particular gate but is assigned to *some* gate under aggregate gate limitations. Constraint (4.15) represents these aggregated constraints. Here, for each aircraft of size $s$ considered in the model (the set of sizes considered in the model reflects the physical gate constraints specific to the airport being modeled), the total number of aircraft at the gate at any point in time cannot exceed $NumGates_s$, the maximum number of size $s$ aircraft simultaneously serviceable at the gates. The final set of constraints (4.16), ensures that at no time are more aircraft at the gate than there are total gates.

## 4.4 Objective Function

The rationale behind the objective function is discussed in Section 2.1. Additionally, the calibration and validation processes that determined the objective function coefficients are discussed in Section 5.2.

The objective function minimizes total delays in passenger-minutes, including gate availability delay, arrival delay, full departure delay[1], and missed passenger connection delays. This leads to the following objective function:

*MINIMIZE:*

$$\begin{aligned} \sum_{ac} GateCost_{ac} \times Delay_{ac} + \\ \sum_{ac} (ArrCost_{ac,t} \times Gate_{ac,t}) + \\ \sum_{ac} (DepCost_{ac,t} \times Dep_{ac,t}) + \end{aligned} \tag{4.17}$$

[1]By using the departure delay multiplier discussed in Section 3.4, the objective function captures the downstream effects of departure delay.

$$\sum_{ac1}\sum_{ac2} NoCnxCost_{ac1,ac2} \times (1 - Cnx_{ac1,ac2})$$

where $SchArr_{ac}$ and $SchDep_{ac}$ represent the scheduled arrival and departure times, respectively, for aircraft $ac$, $GateCost_{ac}$, $ArrCost_{ac}$ and $DepCost_{ac}$ represent the cost coefficients for gate availability delay, arrival delay and departure delay for aircraft $ac$, respectively, and $NoCnxCost_{ac1,ac2}$ represents the cost coefficient of a missed connection between aircraft $ac1$ and aircraft $ac2$. The values of these cost coefficients are discussed in detail in Section 5.2.

## 4.5 Boundary Effect

In a time horizon of 3-4 hours, there are typically some aircraft that depart at the beginning of the time horizon, meaning the aircraft arrive prior to the time horizon considered. There may also be some aircraft that arrive at the end of the time horizon, meaning the aircraft depart after the time horizon considered. This effect is referred to as the "boundary" effect. To account for this boundary effect, each aircraft is considered in one of three sets: $Turn$, which is the set of aircraft that arrive *and* depart in the given time horizon; $NoDep$, which is the set of aircraft that arrive, but do not depart, in the given time horizon; and $NoLand$, which is the set of aircraft that depart, but do not land, in the given time horizon. These sets are further grouped into the sets of arriving aircraft, $ArrPlanes = Turn \bigcup NoDep$, departing aircraft, $DepPlanes = Turn \bigcup NoLand$ and all aircraft by $AllPlanes = Turn \bigcup NoLand \bigcup NoDep$.

These boundary effects determine the set of aircraft over which the variables are defined. For example, it is unnecessary to defined $Land_{ac,t}$ and $Gate_{ac,t}$ for any $t$ and any $ac \in NoLand$. The aircraft indices over which each of the variables is defined is included in Table 4.3.

These boundary effects also determine the form of some of the constraints. For example, we could have an aircraft in the set $NoDep$ which *lands* in the designated time horizon, but doesn't necessarily arrive at a gate before the end of the time horizon. Therefore, Constraint (4.8) is adjusted for the aircraft in the set $NoDep$. Specifically, the constraint corresponding to the last time period, $T$, is omitted. The omission of this constraint changes the definition of $Gate_{ac,T}$ for $ac \in NoDep$ to mean aircraft $ac$ arrives at some gate *at* or *after* time $T$.

The exact form of the constraints in the model adjusting for the boundary effects is included in Appendix A.

## 4.6 Column Reduction

In order to reduce the number of decision variables considered in the model, we restrict the set of time and aircraft indices over which the variables are defined.

As explained in Section 3.1.1, the parameter $Window$ in ASM restricts the possible landing times assigned to an aircraft. In particular, an aircraft can be assigned a landing time within $\pm Window$ minutes from its input, or actual, landing time. Therefore, each aircraft $ac$ has at most $2 \times Window + 1$ feasible landing times. The bounds on this set are known. Namely, for all $ac \in ArrPlanes$

$$\begin{aligned} FirstLand_{ac} &= ActLand_{ac} - Window \\ LastLand_{ac} &= ActLand_{ac} + Window \\ FirstGate_{ac} &= FirstLand_{ac} + Taxi_{ac} \end{aligned}$$

where $ActLand_{ac}$ is the actual landing time for aircraft $ac$. The variables $Land_{ac,t}$ are defined only for $FirstLand_{ac} \leq t \leq LastLand_{ac}$. Furthermore, the variables $Gate_{ac,t}$ are defined only for $t \geq FirstGate_{ac}$.

For schedule integrity, an aircraft is constrained such that it can not push back before scheduled departure time. However, as presented in Section 3.1.2, if an aircraft's actual departure time is before its scheduled departure time, it is constrained such that it can not push back before its *actual* departure time. Additionally, an aircraft can not depart until it has landed, taxied and turned. Denote the earliest time an aircraft can push back as $FirstDep_{ac}$. Its formal definition is:

$$FirstDep_{ac} = \max\left(\min\left(ActDep_{ac}, SchDep_{ac}\right), FirstLand_{ac} + Taxi_{ac} + Turn_{ac}\right)$$

where $ActDep_{ac}$ is the actual departure time and $SchDep_{ac}$ is the scheduled departure time

| Variable | First Index | Second Index |
|---|---|---|
| $Land_{ac,t}$ | $ac \in ArrPlanes$ | $t \in ACLandTimes_{ac}$ |
| $Gate_{ac,t}$ | $ac \in ArrPlanes$ | $t \in Time$ s.t. $t \geq FirstGate_{ac}$ |
| $Dep_{ac,t}$ | $ac \in DepPlanes$ | $t \in Time$ s.t. $t \geq FirstDep_{ac}$ |
| $Delay_{ac}$ | $ac \in ArrPlanes$ | |
| $Connect_{ac1,ac2}$ | $ac1 \in ArrPlanes$ | $ac2 \in DepPlanes$ s.t. $pax^{ac2}_{ac1} > 0$ |
| $NumRes_{ac,t}$ | $ac \in AllPlanes$ | $t \in Time$ s.t. $t \geq FirstGate_{ac}$ |

Table 4.3: Indices over which variables are defined

for aircraft $ac$. The variables $Dep_{ac,t}$ are defined only for $t \geq FirstDep_{ac}$.

The largest contributor to the number of rows in the coefficient matrix is Constraint (4.11). For each time period and for each *pair* of aircraft considered in the model, there is one constraint. In practice, we are only concerned about the connectivity of aircraft if there are passengers that connect between them. If we let $pax^{ac2}_{ac1}$ denote the number of passengers connecting from $ac1$ to $ac2$, we only define $Cnx_{ac1,ac2}$ when $pax^{ac2}_{ac1} > 0$.

The redefined index ranges for the variables are included in Table 4.3.

To quantify the extent to which this re-indexing affects the problem size, four scenarios with varying time horizons were defined and solved using both formulations. Data comparing problem size and solution time are included in Table 4.4. The formulation with restricted indices is referred to as the restricted formulation, denoted Res. in the table, while the formulation with *no* index restrictions is referred to as the full formulation, denoted Full in the table. The four scenarios are based on actual data from January 4, 1998, starting at 16:00 local time. The length of the time horizon is indicated in the table. The table includes the number of aircraft in each set of aircraft types: $Turn, NoLand, NoDep$. The length of the time horizon and the number of aircraft in the model dictate the size of the problem.

The problem is formulated using AMPL version 9.10.27 [8] and solved using CPLEX version 6.6.0 [9]. AMPL has a presolve function that can reduce the number of variables and constraints before sending the problem to CPLEX to be solved. Presolve identifies "fixed" variables, such as constraints (4.4)–(4.6) and substitutes these fixed values everywhere the variables appear in the formulation, eliminating both the fixing constraints and the variables.

| Scenario | 1 | | 2 | | 3 | | 4 | |
|---|---|---|---|---|---|---|---|---|
| Model | Res. | Full | Res. | Full | Res. | Full | Res. | Full |
| Time Horizon (mins) | 90 | | 120 | | 150 | | 180 | |
| $Turn$ | 4 | | 12 | | 16 | | 23 | |
| $NoLand$ | 11 | | 13 | | 14 | | 17 | |
| $NoDep$ | 16 | | 18 | | 22 | | 21 | |
| Original Variables (000) | 4.9 | 11.5 | 8.4 | 21.5 | 13.2 | 32.4 | 18.6 | 45.8 |
| Presolve Variables (000) | 3.6 | 5.8 | 6.3 | 12.0 | 9.9 | 20.7 | 13.9 | 29.8 |
| Original Constraints (000) | 11.0 | 41.1 | 30.2 | 117.5 | 48.5 | 213.2 | 85.2 | 377.0 |
| Presolve Constraints (000) | 3.6 | 12.7 | 11.1 | 78.5 | 27.5 | 147.9 | 33.0 | 246.6 |
| Nonzeros (000) | 12.6 | 105.8 | 64.6 | 324.1 | 150.3 | 644.4 | 272.6 | 1,136.6 |
| Solution Time (secs) | 2 | 6 | 4 | 42 | 7 | 145 | 15 | $\infty$ |

Table 4.4: Comparison of restricted formulation and full formulation

Further, the presolve function tightens bounds on the variables. This tightening could result in some constraints becoming redundant and therefore discarded [6]. The number of variables and constraints before and after applying the presolve function are included in the table. The count of nonzero entries in the matrix is also done after applying the presolve function. The solution times displayed in the table are based on using an 800 MHz Pentium III with 256 MB RAM.

Notice that after the presolve function, there are about twice as many variables in the full formulation and between 3 and 7 times as many constraints, resulting in a coefficient matrix with between 4 and 5 times as many nonzero entries. The larger size of the full formulation significantly impacts the solution time of the model. As shown in the table, the solution time grows exponentially with the problem size. A problem that solved in 15 seconds with the restricted formulation could not be solved with the full formulation due to lack of memory.

# Chapter 5

# Calibration and Validation

There are two major assumptions in the model that were calibrated. The first assumption is the taxi-in time assigned to the aircraft and the second is the objective function coefficients. The first two sections of this chapter explore the calibration process around these two assumptions. The third section provides a validation of the simplifying assumptions regarding gate and ground crew resource availability. Finally, the fourth section compares the results from ASM to the results from a naive model.

## 5.1 Taxi-In Assumptions

As discussed in detail in Section 3.1.3, the taxi-in process is highly stochastic. Further, it was determined that the assignment of accurate taxi-in times to aircraft is not practical in the context of the optimization model. Therefore, various taxi-in assignment schemes were tested and compared, including:

1. assigning a random sample from the unimpeded taxi-in time distribution
2. assigning the average unimpeded taxi-in time (6 minutes)
3. assigning the average overall taxi-in time (7 minutes)

Notice that the resolution of ASM is minutes, therefore all taxi-in times assigned will be an integer number of minutes.

| Date | Sample | | Average Unimpeded | | Overall Average | |
|---|---|---|---|---|---|---|
| | MSE | AVG | MSE | AVG | MSE | AVG |
| Feb 1 | 23.7 | 1.3 | 3.3 | -0.2 | 4.1 | 0.9 |
| Feb 2 | 27.1 | 0.7 | 7.2 | -1.0 | 5.5 | 0.1 |
| Feb 3 | 20.6 | -1.4 | 17.9 | -1.5 | 10.5 | -0.1 |
| Feb 4 | 30.3 | -2.2 | 30.6 | -2.6 | 26.1 | -1.5 |
| Feb 8 | 25.9 | 0.4 | 11.7 | -0.9 | 7.1 | 0.4 |
| Feb 13 | 13.6 | -0.2 | 7.7 | -0.4 | 5.5 | 0.9 |
| Feb 17 | 11.5 | -0.2 | 8.5 | -0.7 | 7.6 | 0.4 |
| Feb 18 | 39.9 | -0.4 | 26.4 | -1.5 | 24.5 | -0.4 |
| Feb 19 | 38.3 | 0.4 | 20.9 | -0.6 | 17.9 | 0.6 |
| Feb 20 | 19.5 | 0.1 | 5.4 | -0.8 | 4.2 | 0.4 |
| Feb 24 | 18.0 | -0.2 | 7.3 | -0.7 | 7.1 | 0.4 |
| Feb 27 | 34.6 | 0.3 | 13.9 | -1.2 | 10.7 | -0.1 |
| Average | 25.3 | -.01 | 13.4 | -1.0 | 10.9 | 0.2 |
| Standard Deviation | 9.2 | 0.9 | 8.7 | 0.6 | 7.7 | 0.7 |
| Lower Bound | 20.0 | -0.6 | 8.5 | -1.4 | 6.5 | -0.2 |
| Lower Upper | 30.5 | 0.4 | 18.4 | -0.7 | 15.3 | 0.5 |

Table 5.1: Comparison of taxi-in assumptions

Since the arrival time estimates in a particular scenario are interdependent (the aircraft share finite gate resources), multiple independent scenarios were considered. In each scenario we set the parameter $Window = 0$, meaning that the landing times are assumed fixed. Therefore, ASM is, in effect, predicting the arrival at gate and push back times. The metric compared in this analysis is the arrival error, defined as the difference between ASM's predicted arrival time and the actual arrival time. For each scenario and each assignment scheme, the average arrival error and the average mean-squared arrival error were calculated, with the results included in Table 5.1. The data included in the analysis are for twelve days in February 1998 from 16:00 to 19:00 local time.

Notice that the results from the overall average have the lowest mean-squared arrival error. Further, the 95% confidence interval of the average arrival error covers zero. Therefore, we assume a constant taxi-in time set to the overall average of the taxi-in times at Airport 1, or 7 minutes.

## 5.2 Objective Function Assumptions

Recall that the objective function, as given in (4.18), is designed to minimize the total of four categories of delay: gate availability delay, arrival delay, departure delay and delay incurred by passengers who miss their connecting flight. The objective function is measured in minutes of passenger delay. In finding a solution, ASM makes trade-offs between these delays. For example, ASM decides whether to delay a departing aircraft for delayed incoming passengers. In making this decision, ASM considers the relative costs of departure delay and missed connection delay. Therefore, it is important that these relative costs are accurately reflected in the model. This section describes the process of calibrating and validating the cost coefficient data.

The cost coefficients were calibrated on a set of 12 independent time horizons from 12 days in February 1998. The time horizons ranged from three to four hours in length. The cost coefficients were tuned on these data and the resulting coefficients were validated on another set of 12 independent time horizons from 12 days in January 1998. For this analysis, the landing times are assumed fixed; ASM is not allowed to resequence the arriving aircraft. This assumption means that ASM is essentially predicting arrival at gate and departure times. The judgement of the accuracy of the cost coefficients is based on comparing ASM predicted to actual movement times.

First, the cost coefficient for gate availability delay was tested. Recall that the variable $Delay$ measures the minutes of delay incurred while an aircraft waits for a gate to become available. Therefore, the coefficient is a multiplier of the minutes of delay that translates the delay into passenger-minutes of delay. The coefficient $GateCost_{ac}$ was tested at three different values:

**number of passengers** This converts aircraft minutes of delay into passenger-minutes of delay

**zero** This can be interpreted to mean that gate availability delay is incorporated in arrival delay

**small positive** This tests the sensitivity of the solution to the coefficient value

The results of the calibration study show that if the gate availability delay cost is set to zero, ASM produces a solution in which each aircraft arrives at the gates at its scheduled arrival time or its first possible arrival time, whichever is later. In other words, the ASM solution contains no early arrivals, which yields an arrival delay distribution significantly different than the actual. From these results, we concluded that the gate availability delay coefficient should be positive.

For the 12 independent time horizons tested, there is no difference in the solutions generated with the coefficient equal to the number of passengers and with the coefficient equal to a small positive value. This is because for these 12 scenarios, the gate availability constraints are never binding, meaning there is always at least one gate available. Given this result, we set the gate availability delay coefficient to be the number of passengers on board the aircraft, i.e.,

$$GateCost_{ac} = pax_{ac}^{a}.$$

where $pax_{ac}$ indicates the number of passengers arriving on aircraft *ac*. This coefficient is the multiplier that translates minutes of aircraft delay into passenger minutes of delay.

The arrival delay cost coefficient is multiplied by a binary variable, meaning the coefficient is measured in passenger minutes of delay. Further, the coefficient is really a *function* over time. Since we know that the gate constraints are not binding in these 12 scenarios, we know that changing the arrival delay cost coefficient will not alter the results. Therefore, we use intuition in determining the cost coefficient. There is little operational benefit from an early arrival; the schedule is designed with adequate time for turning plus a little slack. In some cases, an early arrival could usurp resources from other aircraft or an early arrival might have to wait for a gate to become available. Given these conditions, we do not want a cost function that encourages early arrivals. Then, for instances of positive arrival delay, the cost is set to total passenger minutes of delay, i.e. the product of the minutes the aircraft arrives after scheduled arrival time and the number of passengers onboard the aircraft. The resulting function is:

$$ArrCost_{ac,t} = \begin{cases} 0 & \text{if } t < SchArr_{ac} \\ pax^a_{ac} \times (t - SchArr_{ac}) & \text{otherwise} \end{cases}$$

The missed passenger connection coefficient, $NoCnxCost_{ac1,ac2}$, is multiplied by binary variable $Cnx_{ac1,ac2}$ in the objective function. It is set so that this objective function term represents minutes of passenger delay due to misconnections. The delay incurred by each passenger is the time from the originally scheduled departure to the next departure to the same destination. Therefore, the missed passenger connection delay coefficient is set to the product of the time until the next departure to the same destination, denoted $Next_{ac}$, and the number of connecting passengers, i.e.

$$NoCnxCost_{ac1,ac2} = Next_{ac2} \times pax^{ac2}_{ac1}.$$

The departure delay coefficient, $DepCost_{ac,t}$, is multiplied by binary variable $Dep_{ac,t}$. $DepCost_{ac,t}$ is determined so that this product represents minutes of passenger delay. Like the arrival cost coefficient, the departure cost coefficient will be a function over time. In each of the 12 scenarios tested, there was some trade-off between missed connection delay and departure delay. Therefore, a more detailed analysis was done to calibrate this cost function.

Since the passenger connection data used in ASM is simulated rather than observed data, the departure time decisions made for a particular aircraft are likely to deviate from the actual. This observation means that a comparison of departure errors by aircraft is not meaningful. Therefore, the coefficients were chosen by comparing the *distributions* of actual departure delay and the departure delay predicted by ASM.

In this analysis, we consider three categories of departure delay: negative delay (early departure), "slight" delay and "significant" delay. The cost of an early departure can either be zero, which means there is no benefit to departing early, or it can be negative, which means there is some benefit to departing early. If a negative cost is assumed, it is set to be the negative of the number of minutes the aircraft left before its scheduled departure time. For example, if the scheduled departure time is 10:10 and the aircraft departs at 10:07, the departure cost corresponding to 10:07 is $-(10:00 - 10:07) = -3$. Due to the slack built into the system, an aircraft can often recover from a slight delay at departure. Therefore, a slight delay is assumed to have a different cost function. In these scenarios, a slight delay is interpreted to be either 10 or 15 minutes. The cost of the slight delay is either zero or a linear cost function representing the Direct Operating Costs (DOC), as discussed in Section 3.4. Finally, significant delays are defined as any delay that is neither negative nor slight. In these scenarios, the base cost for a significant delay is the product of the number of passengers onboard the aircraft and the minutes of departure delay. The base cost for the significant delay is then either multiplied by the Delay Multiplier (DM), as described in Section 3.4, or it is not. In this analysis the following four scenarios were tested:

1. Early cost is zero, slight cost is 0 for 10 minutes, significant cost includes DM
2. Early cost is negative, slight cost is DOC for 15 minutes, significant cost includes DM
3. Early cost is negative, slight cost is DOC for 10 minutes, significant cost includes DM
4. Early cost is negative, slight cost is DOC for 10 minutes, significant cost excludes DM

The delay distributions corresponding to these scenarios are included in Figure 5-1. As shown in the first chart in the figure, scenario 1, in which departure delays up to 10 minutes have zero cost, results in a delay distribution that is significantly different than actual. In particular, most (75%) of the aircraft depart at the last no-cost time period, or 9 minutes after scheduled departure. This result implies that a continuous, rather than a step, cost function is necessary. Scenarios 2 and 3 are examples of such continuous cost functions. In these two scenarios, the early departures receive a slight benefit (set to the negative of the minutes the aircraft depart early) and the departure delays up to 15 and 10 minutes, respectively, have a linear cost function based on the direct operating costs. The resulting delay distributions are included in the second and third charts of Figure 5-1. Notice that having the linear cost function for up to 15 minutes of delay in scenario 2 results in a significant deviation in the number of aircraft departing with 12 minutes of delay. However, this deviation is not apparent in the results from scenario 3. This comparison indicates that the cost function in scenario 3 yields a solution that more accurately reflects observed behavior. The final scenario considers the impact of the delay multiplier on the movement time decisions. The only difference between scenarios 3 and 4 is that scenario 4 does not use the delay multiplier to estimate downstream delays. As shown in the fourth chart in Figure 5-1, the exclusion of the delay multiplier in the cost function has negligible impact on the solution.

Based on these results, the third cost function is used in ASM[1]. To validate this choice, the cost function was tested for 12 days in January 1998 from 16:00–19:15 local time. The resulting distribution of delay is included in Figure 5-2. Notice that the delay distributions for January are very similar to those for February.

It is important to note that the data set used to generate Figure 5-2 excludes aircraft for which the difference between actual departure delay and actual arrival delay exceeds 40 minutes. These data points were excluded because it is unlikely that delays of that magnitude (greater than 40 minutes) can be explained by ground crew resources availability, gate availability or passenger connections. Therefore some factor(s) external to ASM influenced the departure time. Despite the omission of these identified points, ASM still significantly underestimates the number of departures incurring delays exceeding 40 minutes. However, these excessive observed delays are still likely attributable to factors external to ASM. For example, if an aircraft is observed to have arrived 15 minutes *early* but departed 25 minutes late, the corresponding data point is included in the chart. However, it is unlikely that the ASM solution will delay an aircraft by 25 minutes based on the factors considered in the model.

The most significant difference in the distributions exists for departure delays of 1–4 minutes. ASM assigns on-time departures to aircraft that were actually delayed 1–4 minutes. The departure process is an extremely complex process involving the synchronization of many resources and sub-processes. Before an aircraft is ready for departure the passengers must deplane the arrival and board the departure, baggage handlers must unload and load the baggage, caterers and cleaning crews must remove rubbish and replenish food and beverage

[1]To confirm that the use of the delay multiplier does not affect the solution, the results discussed in Section 6 are calculated both including and excluding the delay multiplier; the delay multiplier has negligible impact.

supplies, the aircraft must be checked for flight safety and refueled, the cabin and cockpit crews must arrive and prepare for departure and so forth. Variability exists in each of these sub-processes. A delay of a few minutes could be caused by numerous factors external to ASM.

Note that the approach to this calibration process is subjective in that the decision about the parameters are based on visual interpretations. Additionally, it considers only a few of the possible cost functions. It could be possible to further adjust the objective function coefficients in ASM to produce solutions which more accurately reflect observed behavior. However, optimizing these parameters is an extremely difficult challenge. It is impossible to know *a priori* how a modification of the cost coefficients will affect the solution. Furthermore, and more importantly, it is difficult to identify an optimal solution. It is unknown how much of the deviation from actual departure times is attributable to the use of sub-optimal parameters and how much is attributable to including insufficient information about the turning process in the model. For the purposes of this study, the level of detail of this analysis is adequate.

## 5.3 Disaggregation of Results

As discussed in Section 2.3, the solutions generated by ASM are not necessarily feasible. In order to manage the number of variables and constraints in the model, some aggregated constraints are used. In particular, aircraft are not assigned to specific gates. Instead, they are assigned to *some* gate in the set of gates considering aggregate constraints.

For example, consider an airline with three gates (1,2,3) and three aircraft types (A,B,C). Gate 1 can service all three aircraft types but gates 2 and 3 can service only aircraft type C. When aggregated, these conditions are modeled in ASM as no more than one type A, one type B, three type C and three total aircraft can be at the gates at the same time. Given these constraints, ASM's solution could assign a type A and a type B aircraft at the gates at the same time, for which no feasible gate assignment exists.

Further, baggage handlers are not assigned to aircraft in teams. Nor are they assigned

| Aircraft | First Time At Gate | Last Time At Gate | Resources Assigned in Time Units (1,2,3,4) |
|---|---|---|---|
| A | 1 | 4 | (2,0,0,2) |
| B | 2 | 3 | (0,2,0,0) |
| A | 3 | 3 | (0,0,2,0) |

Table 5.2: Possible infeasible solution from ASM

continuously to aircraft. Instead, in each time unit baggage handlers are assigned to an aircraft independent of their previous assignment. Given this, ASM could produce a solution that is not feasible.

For example, ASM could produce the solution in Table 5.2 for three aircraft A,B,C over a time horizon of four time units (1,2,3,4). Notice that with this solution, the assignment of baggage handlers to aircraft A is not continuous. In particular, there are no baggage handlers assigned during time units 2 and 3. However, there is no way to form a continuous assignment of resources to all three aircraft without either reducing the number of resources needed in the time period (in effect, changing the arrival and/or departure time from the gate) or assigning more than 2 baggage handlers at the same time. Assuming that only 2 baggage handlers are available, it is obvious that the above solution produced by ASM is not feasible.

In addition, ASM does not consider any contractual issues regarding baggage handlers. For example, it does not explicitly consider scheduled breaks for the crew. Instead, the number of baggage handlers input into ASM is reduced from the actual number available. For the analyses in this study, the number of baggage handlers assumed in the model is 60, while the airline schedules 80 during peak hours[2] This assumption implies that there is slack designed in ASM's solution regarding baggage handlers. Therefore, given the full set of baggage handlers, a feasible assignment of baggage handler teams to aircraft considering contractual issues should exist.

[2]Very few of the analyses included aircraft delayed due to a shortage of resources when the number of resources available is assumed 60. Furthermore, such delays had a magnitude of a few minutes. Therefore, tightening the constraint on baggage handler availability is not expected to affect the solution significantly.

To determine the limits of these assumptions, a second model was built that assigns aircraft to gates and teams of baggage handlers to aircraft *given* the aircraft movement times from the ASM solution. The formulation of this model can be found in Appendix B. This model is intended to justify the aggregation assumptions of ASM. It assigns each aircraft to a specific gate and assigns to it a team of baggage handlers from the set of baggage handlers who are assigned to the zone that includes that gate. The model is designed such that it counts the number of *additional* resources required in order to achieve the given movement times. If the minimum number of additional resources required is zero, then we know the ASM solution is feasible. If, however, the minimum number of additional resources required is positive, then the solution is infeasible and the number of additional resources required indicates the extent of the infeasibility. If, on the other hand, the problem is *infeasible*, then we know the ASM solution is infeasible because no feasible gate assignment exists.

This feasibility model is not efficient. Therefore, it was not run on every ASM solution. However, a number of solutions were tested and each tested solution was found to be feasible. All ASM solutions from which the calibration and validation studies are drawn were tested for feasibility. The result is that a feasible gate and baggage handler team assignment could be found for each of the solutions.

## 5.4 Comparison to Naive Model

To determine whether ASM is effective in predicting arrival and departure times, its solutions, assuming again that the parameter $Window = 0$, were compared to that of a naive model. The naive model is designed with a constant taxi-in time set to the average of the overall taxi-in time and with a constant turn time based on the minimum scheduled turn time. The naive model has the same taxi-in assumption as ASM. The difference between the models is that ASM considers ground crew resources and passenger flows in determining departure time.

Recall from Section 5.1 that the gate constraints are not binding in any of the scenarios tested to date. As a result, the arrival time estimates of ASM and the naive model are

| Scenario | ASM Departure Error | | Naive Model Departure Error | |
|---|---|---|---|---|
| | MSE | Avg | MSE | Avg |
| 1 | 168.66 | -0.33 | 882.21 | 6.28 |
| 2 | 46.95 | -2.29 | 113.31 | 0.98 |
| 3 | 46.00 | -0.08 | 107.54 | 3.74 |
| 4 | 33.04 | -1.71 | 157.02 | 3.95 |
| 5 | 697.38 | -5.27 | 1,347.09 | -3.42 |
| 6 | 563.72 | -9.08 | 1,160.26 | -5.86 |
| 7 | 288.54 | -3.73 | 646.04 | 1.88 |
| 8 | 12.67 | -0.67 | 53.27 | 1.78 |
| 9 | 16.78 | 1.07 | 396.37 | 8.22 |
| 10 | 122.09 | -5.00 | 382.53 | 2.98 |
| 11 | 572.92 | -7.81 | 2,409.85 | 0.69 |
| 12 | 113.02 | -4.47 | 1,757.24 | 6.07 |
| Average | 223.48 | -3.28 | 784.39 | 2.28 |
| Std. Error | 248.32 | 3.17 | 752.20 | 3.98 |
| Lower Bound | 82.98 | -5.08 | 358.81 | 0.02 |
| Upper Bound | 363.98 | -1.48 | 1,209.98 | 4.53 |

Table 5.3: Comparison of departure error between ASM and Naive Model

identical. The comparison presented here therefore focuses on the difference of the departure time predictions.

Since the departure times of the aircraft in a particular scenario are interdependent in ASM (the aircraft share finite ground crew resources), multiple independent scenarios were considered in order to compare the two models. The metric considered is departure error, defined as the model's prediction of departure time minus the actual departure time. For each scenario and for each model, the average departure error and the mean-squared departure error were calculated. The results are included in Table 5.3. The data included in the analysis are for twelve days in January 1998 from 16:00–19:15 local time.

Notice that the mean-squared departure errors for the two models are significantly different; the errors from ASM are generally significantly smaller than the naive model's errors. In fact, a Wilcoxon signed rank test confirms that the MSE values for ASM are less than those for the naive model with a two-tailed significance level of 0.2%. This implies that the additional factors considered in ASM are influencing the turn process and are improving the departure time predictions. However, the confidence interval of the average departure delay for ASM does not cover zero. In fact, the confidence interval consists of an entirely negative range. This means that ASM's departure time estimates tend to be earlier than the actual departure time, implying there is some bias in the predictions.

This result, however, is consistent with the results from the analysis of the cost function in Section 5.2. In that section, we found that ASM's departure predictions underestimate delays of 1–4 minutes and delay exceeding 40 minutes. By underestimating these delays, the average delay will tend to be negative.

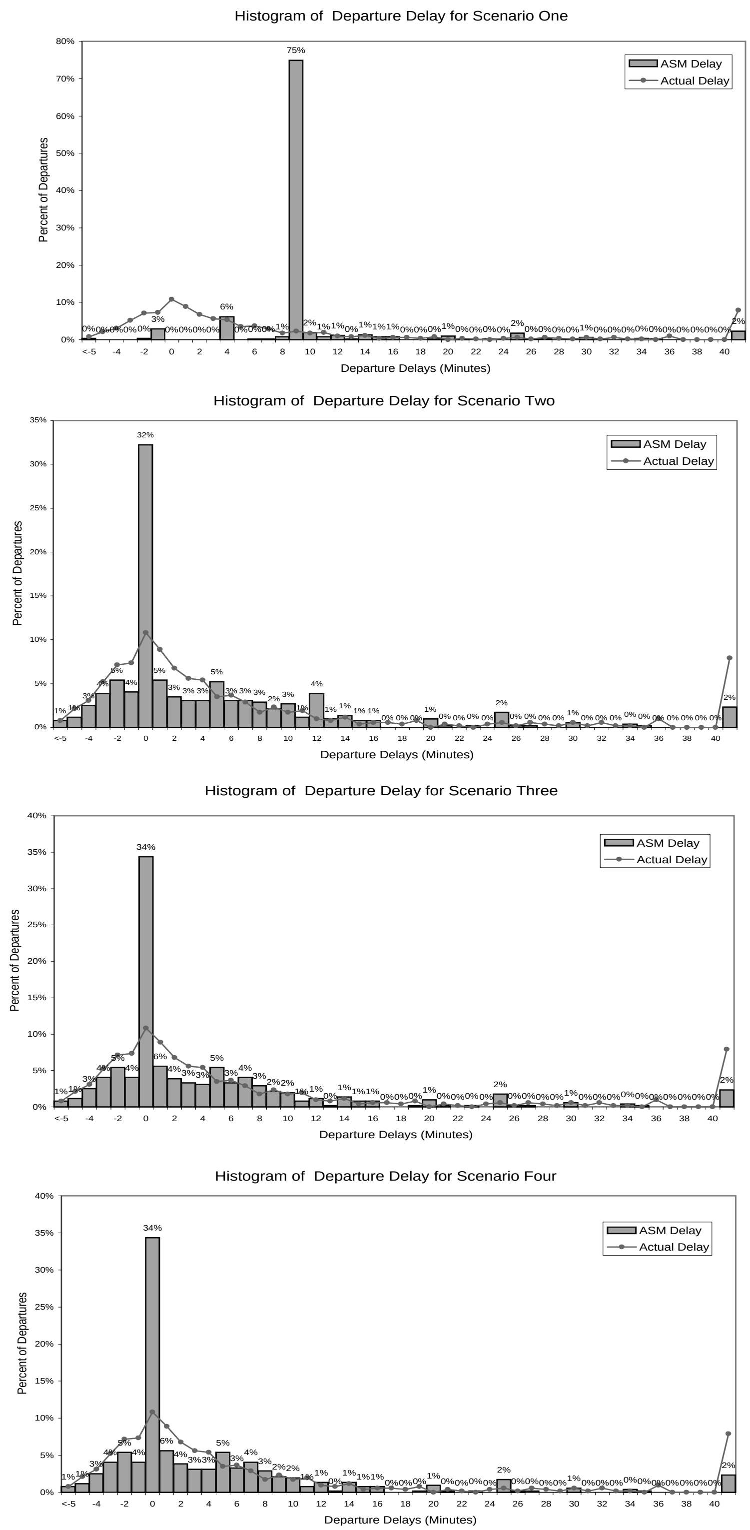


Figure 5-1: Delay distributions assuming four different departure cost functions

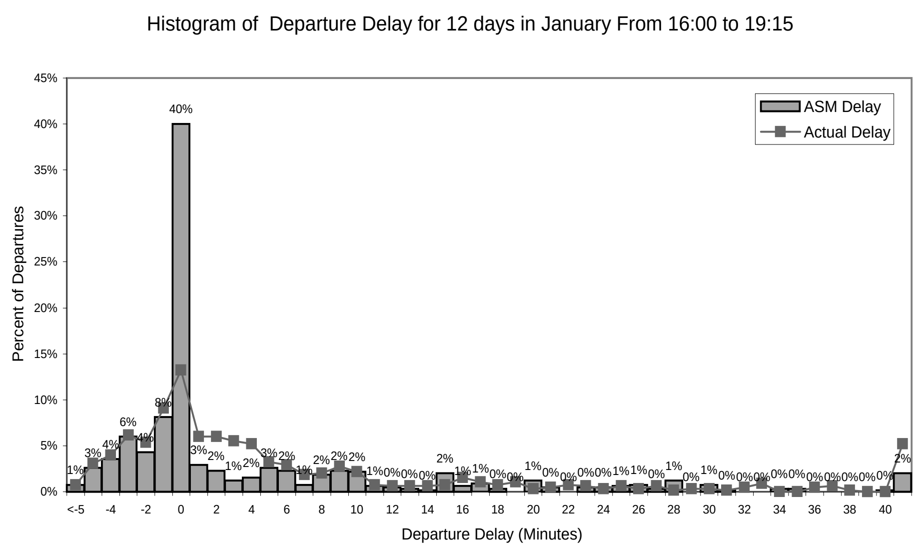


Figure 5-2: Validation of cost coefficients for 12 days in January 1998

# Chapter 6

# Results

This section explores the results of the studies conducted using the Arrival Sequencing Model. Recall that the objective of this study is two-fold: to estimate the benefits of more accurate arrival time predictions and to estimate the benefits of sequencing flexibility in the arrival process. To obtain such estimates, scenario analyses were conducted using ASM. A description of the approach and a discussion of the results to date are included here.

Some of the analyses included in this chapter consider two days in January 1998; both days cover the same time period, namely 16:00–19:15 local time. These days were chosen because they represent two very different operating conditions for the airline. On one day, 84% of the aircraft arrive and depart at most 15 minutes after the scheduled arrival and departure times. Furthermore, no aircraft experienced either an arrival delay or a departure delay exceeding 60 minutes. This day is referred to as the "On-Time" day. Flights operating on the other day experience much more delay. Only 64% of the aircraft arrive and depart within 15 minutes after the scheduled arrival and departure times. Furthermore, 14% of the aircraft experienced either an arrival delay or a departure delay exceeding 60 minutes. This day is referred to as the "Busy" day. Figure 6-1 plots the arrival and departure delay for each aircraft for the two days.

In addition to significantly different observed delays, the two days have different passenger connection statistics. In solving the On-Time day problem, ASM decides to delay only two out of 56 departing aircraft to accommodate connecting passengers. On the other hand,

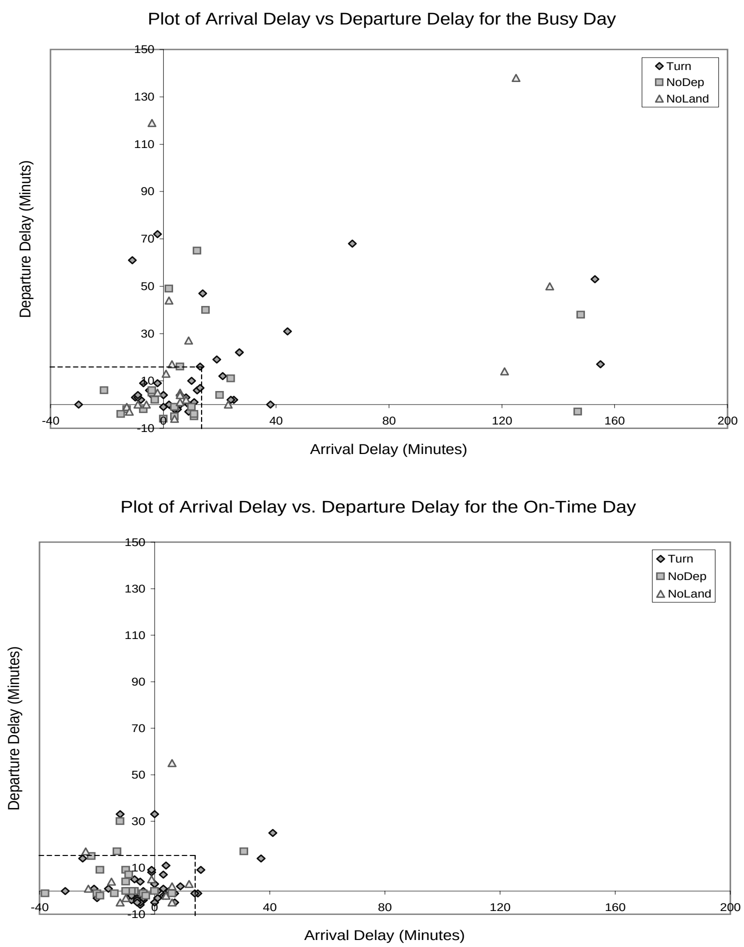


Figure 6-1: Delays for time horizons considered in the results analyses

ASM decides to delay 28 out of 61 departing aircraft to accommodate connecting passengers on the Busy day.

## 6.1 Benefits of More Accurate Arrival Time Predictions

Recall that one of the inputs to ASM is the estimated landing time for each aircraft. Assuming these landing times, ASM optimally allocates resources to the aircraft at the gate and makes decisions about whether to hold departing aircraft for incoming passenger connections. However, as shown in Figure 6-2, these estimates are not perfect. Therefore, the airlines are making these hold/no-hold decisions under uncertainty.

In order to measure the impact of this uncertainty, we need to understand how changes in the landing time estimates affect decision making. In particular, we want to understand how this uncertainty affects the quality of the decision to hold departing aircraft for passengers delayed on incoming flights. To do this, we compare the passenger minutes of delay incurred under perfect information with the passenger minutes of delay incurred when the landing time estimates are uncertain. The difference indicates the passenger minutes of delay that could be saved with more accurate landing time estimates.

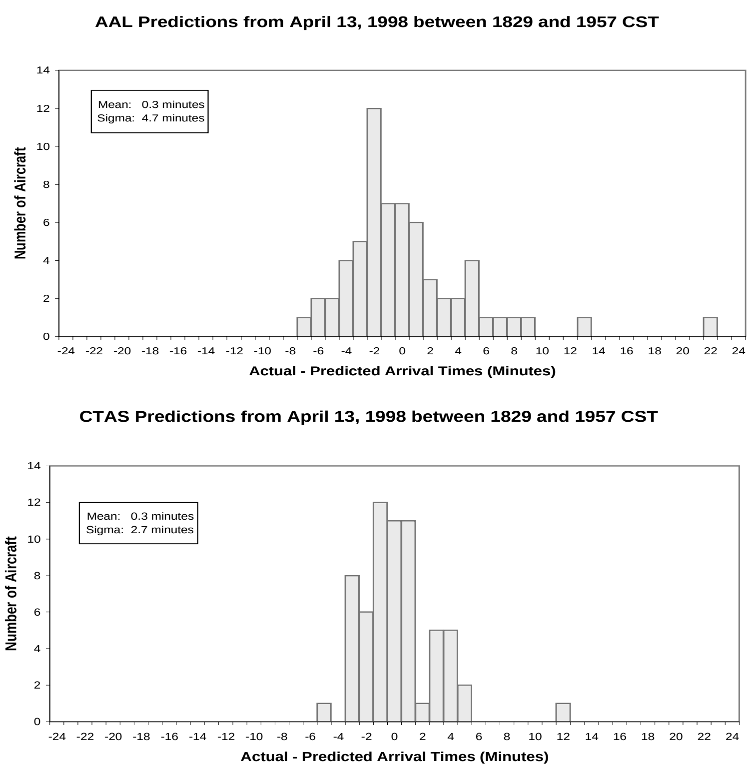


Figure 6-2: Comparison of CTAS arrival time estimates are those currently used by American Airlines at DFW

Obtaining this difference is a multi-step process. Figure 6-3 illustrates the steps in the process. First, to represent the case of perfect information, the actual landing time data are input into ASM and the corresponding objective function value is captured. This corresponding objective function value is referred to as the *base case value*. Then, the landing time estimates are perturbed; a sample from a normal distribution with mean zero and a specified variance is added to each landing time. These landing times represent the estimated landing times. The optimal solution generated by ASM based on the estimates is used as the airline's plan of operations. Recall that the solution generated by ASM includes variables $Cnx_{ac1,ac2}$, which indicate whether passengers from aircraft $ac1$ are able to connect to aircraft $ac2$. These variables represent the hold/no-hold decisions made by the airline based on the noisy estimates.

These hold/no-hold decisions must be made in advance for a tight connection to be made successfully. To understand how the execution of these decisions is affected by the uncertainty, the values of the variables are temporarily fixed in ASM, meaning constraints representing the value of the $Cnx$ variables are added to the model. Then this "augmented" model is solved using the base case data to represent the execution. The objective function value of this augmented problem is referred to as the *perturbed value*. The difference between the base case value and the perturbed value indicates the impact of uncertainty on the decisions.

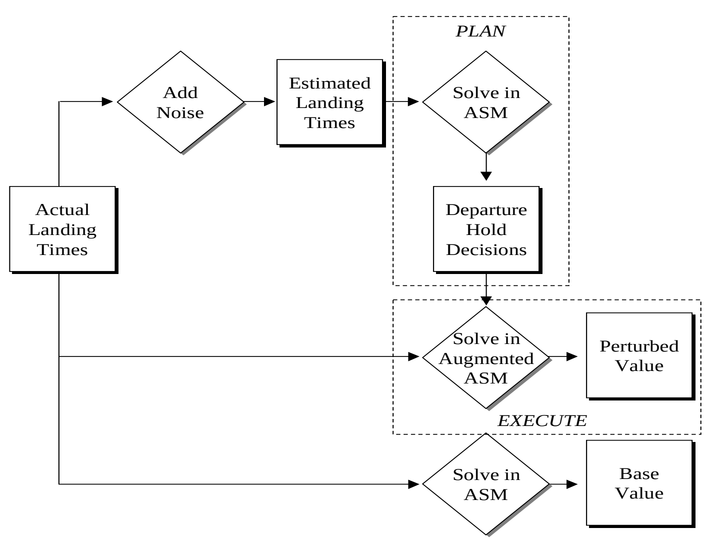


Figure 6-3: Approach to measuring benefit of accurate landing time estimates

Figure 6-2 indicates that the standard deviation of the current arrival time estimates is around 5 minutes. However, other data indicate that the standard deviation of the current estimates can be as great as 7 minutes. CTAS, on the other hand, reports estimates with a standard deviation of 3 minutes. In order to determine how these varying levels of uncertainty affect decision making, we used the approach outlined above.

Recall that the noise added to the actual landing times is a sample from a normal distribution centered at zero with a given standard deviation. A range of standard deviations was considered, from 2 to 7 minutes. For each standard deviation, 12 distinct perturbed data sets were created, yielding 12 perturbed values. The percentage difference between the perturbed values and the base case value were captured. The average and the standard deviation of these differences are captured in Figure 6-4 for the Busy day and the On-Time day.

There are a number of important observations from Figure 6-4. First, the percent increase in passenger minutes of delay increases steadily on the Busy day once the standard deviation of the noise exceeds 3 minutes. This result indicates that during time periods where many passengers are making connections under delayed conditions, increasing the accuracy of the landing time estimates increases operational efficiency as measured in passenger minutes of delay incurred. On the other hand, the relationship between landing time estimate accuracy and operational efficiency is not as clear for the On-Time day. It is important to note, however, that there is significant variation in the averages presented assuming a 5- and a 6-minute standard deviation. This variation is caused by a few samples from the tails of the normal distribution. If these outlier data are excluded, the results are as shown in Figure 6-5. However, even adjusting for the outlier data, the relationship is much less clear.

The second important observation is that the effect of the noise is greater on the Busy day, regardless of the distribution of the noise. This result indicates that days experiencing more delays and more tight passenger connections are more sensitive to the uncertainties.

Furthermore, since the Busy day has 28 aircraft with over 60 minutes of arrival or departure delay, the objective function value is going to be significantly greater than that for the On-Time day. In fact, the difference between the base case value for the Busy day is 4.5

times that for the On-Time day. These differences will skew the percentage improvement numbers. If we examine directly the number of minutes of incurred delay, the difference between the days is greater, as shown in Figure 6-6. For example, consider the difference between the CTAS estimates and those currently used by American Airlines in managing its ground operations. Recall the CTAS variance is around 3% and the AA variance is about 5%. This difference equates to 2000 minutes of passenger delay during a period of 3.25 hours on the Busy day and 500 minutes of passenger delay during a period of 3.25 hours on the On-Time day. These results indicate that there is a significant improvement in ground operations efficiency when accurate landing time predictions are available, especially for days with significant delay and passenger connectivity.

Recall in Section 5.2 that the results from ASM were negligibly affected by considering the delay multiplier for delays greater than 10 minutes. Further, the cost function assumed in the model considers the delay multiplier. The above analyses were conducted twice, once assuming an objective function that included the delay multiplier and once assuming an objective function that did not. In all other aspects, the two problems were identical. As can be seen in Figure 6-7, the difference in the departure delay cost function has little impact on the results.

The results from these analyses indicate that more accurate landing time estimates increase operational efficiency. The extent of these improvements will vary by the delay conditions of the day and the volume of passenger connections being managed.

## 6.2 Benefits of Influencing Arrival Sequence

Recall that the parameter $Window$ determines the flexibility of the sequencing decisions. In particular, $Window$ determines the number of minutes an airline can move a particular aircraft's landing time up or back within its set of feasible landing times. Increasing the value of $Window$ creates more options for the airline to manage its delay. This section explores the extent to which these additional options improve the operational efficiency.

As can be seen in Figure 6-8, the number of passenger minutes saved increased by about

1–2% for every one-minute increase in the parameter $Window$. Notice that the benefit of increased flexibility appears to be greater for the On-Time day than the Busy day. Recall, however, that 14% of aircraft on the Busy day incurred arrival or departure delays exceeding 60 minutes, compared to zero aircraft with such excessive delays on the On-Time day. This observation implies that an analysis considering the percentage change in the objective function will be somewhat skewed. Therefore, Figure 6-9 shows the number of passenger minutes saved as a function of the parameter $Window$. Notice here that the number of minutes saved increases by about 100 minutes for each unit increase in the parameter $Window$ for the two days until $Window = 11$. For the Busy day, the potential savings when $Window = 12$ increases by almost 10,000 passenger minutes and for $Window > 12$ the potential savings increases at a rate of 100 minutes per unit increase in $Window$. For the On-Time day, the rate of increase slows to about 50 passenger minutes per unit increase in $Window$.

These results indicate that the passenger minutes of delay saved increases as the flexibility of the sequencing increases. Furthermore, both the Busy and the On-Time days see significant savings in delays when preferential sequencing is allowed. Notice that the magnitude of the delay reduction is much more significant from preferential sequencing than from more accurate arrivals times. Recall that 2000 passenger minutes are saved when the variation of the landing time estimates is reduced from 5% to 3% on the Busy day. Assuming that the airlines have flexibility to move aircraft landing times within 5 minutes of the actual ($Window = 5$), the potential savings are between 3000 and 5000 passenger minutes of delay in a 3.25 hour time period.

Again we checked the impact of the inclusion of the delay multiplier in the cost function for departure delay on these results. As indicted in Figure 6-10, considering the delay multiplier does not significantly impact the results. Therefore, we can be comfortable with our departure cost function assumption.

These results indicate that the benefits of preferential sequencing can be very significant. However, it is difficult to extrapolate the potential benefits for a month or a year from a study that considers only two 3.25 hour time periods. Therefore, a study considering a month-long span of data was conducted to estimate the benefits over a longer horizon. In this study, the benefits from February 1998 were estimated. Each day in February was represented in five 3-hour scenarios: 8-11 AM, 11-2 PM, 2-5 PM, 5-8 PM, 8-11 PM. Because this study considers various weather and traffic conditions, the results give a more accurate picture of potential benefits over a longer time horizon.

The average benefits were estimated over a three-hour period. For this analysis, the minutes of passenger delay saved for each of the 5 time horizons over the 28 days (140 scenarios) were estimated for six different $Window$ parameter values. For each $Window$ value, the average of the 140 scenarios was calculated. The results are found in Figure 6-11.

Notice almost 4,800 minutes of passenger delay could be saved in a 3-hour period on average assuming an airline could adjust an aircraft's landing time up to three minutes. Further, an increase of the $Window$ parameter by one minute implies an additional savings of about 600 passenger minutes of delay. In particular, an average of 12,600 passenger minutes of delay could be saved every three hours assuming an airline has 15 minutes of sequencing flexibility.

To better understand these averages, Figure 6-12 shows the distributions of the potential savings by $Window$ parameter value. Notice that as the value of $Window$ increases, the mean of the distribution shifts to the right, as expected from the results stated above, and the distribution widens.

The cumulative potential savings over the entire month at IAH can be seen in Figure 6-13. Notice that allowing an airline 3 minutes of sequencing flexibility can save almost 700,000 passenger minutes of delay in one month or over 8 million passenger minutes of delay in one year. Additionally, an increase in the sequencing flexibility of one minute can result in an increase in savings of 100,000 passenger minutes of delay per month or 1.2 million passenger minutes of delay per year. Assuming 15 minutes of sequencing flexibility, an airline could save 1.8 million passenger-minutes of delay per month or 21.6 million passenger minutes of delay per year.

Recall that the potential savings cited in this chapter are for *one* airline at *one* hub airport. These results can extend to at least 15 hub airports. Therefore, these results imply incredible potential savings in delay due to increased collaboration during the arrival process.

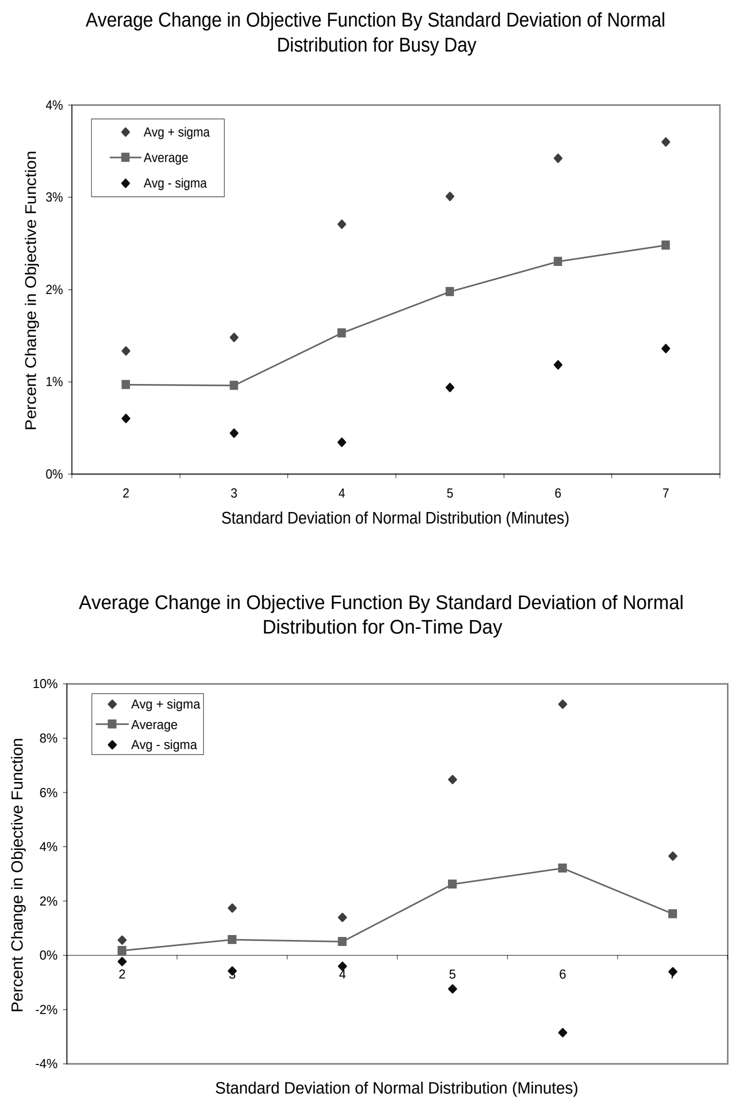


Figure 6-4: Impact of uncertainty in landing time estimates on operational efficiency

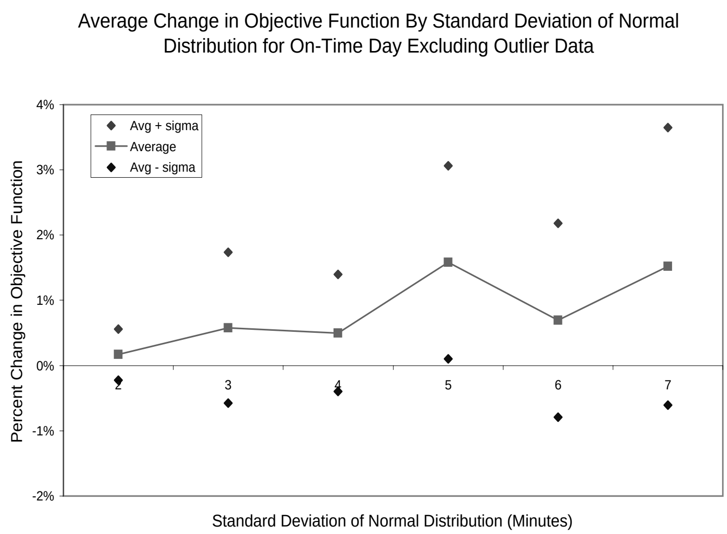


Figure 6-5: Impact of uncertainty in landing time estimates on operational efficiency on January 2, 1998 excluding outliers

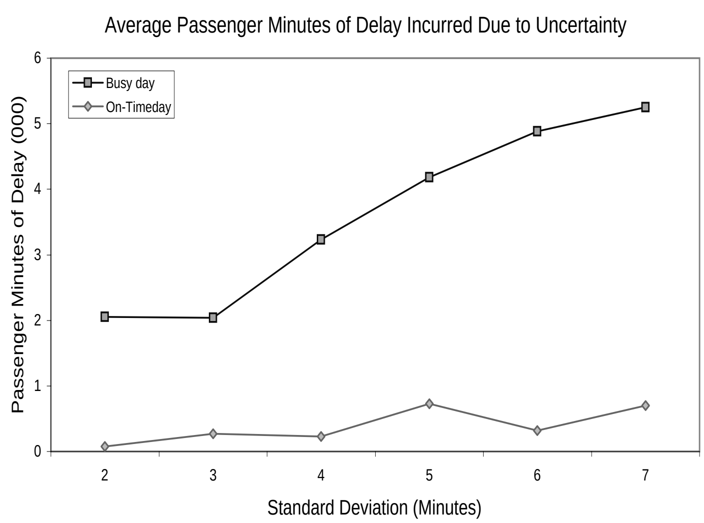


Figure 6-6: Passenger minutes of delay incurred by making decisions under increased uncertainty

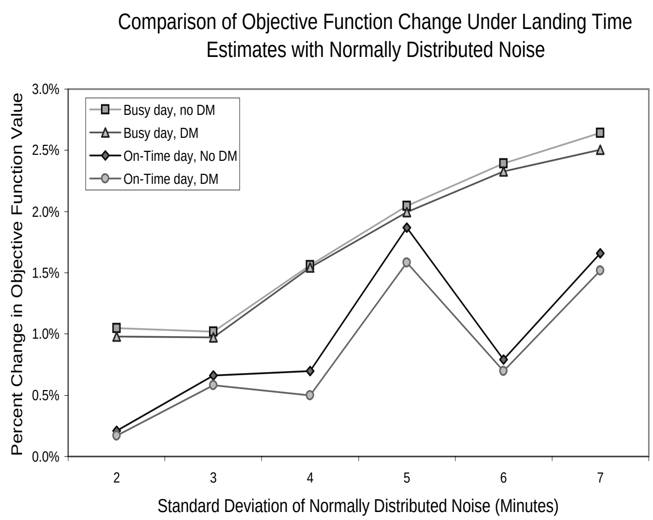


Figure 6-7: Effect of the delay multiplier on results regarding arrival time estimates

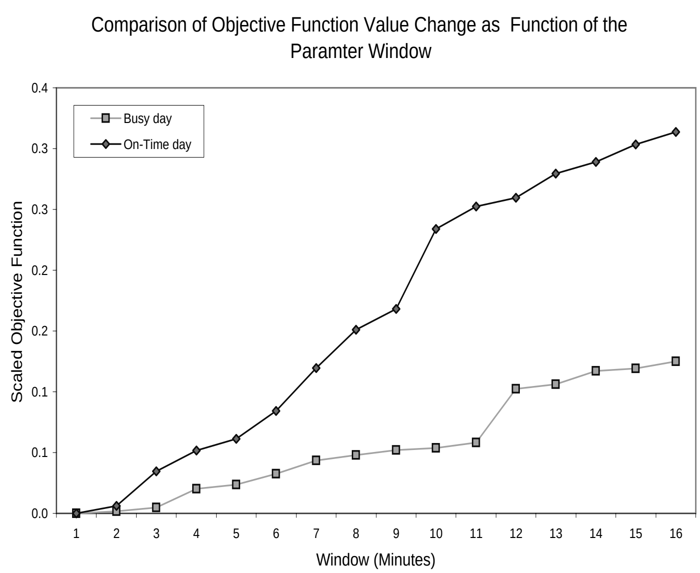


Figure 6-8: Impact of increased flexibility in determining the arrival sequence on operational efficiency

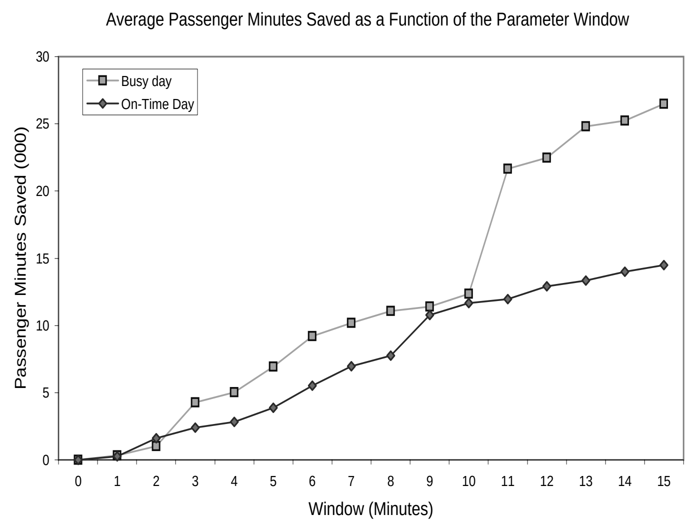


Figure 6-9: Passenger minutes saved as a function of sequencing flexibility

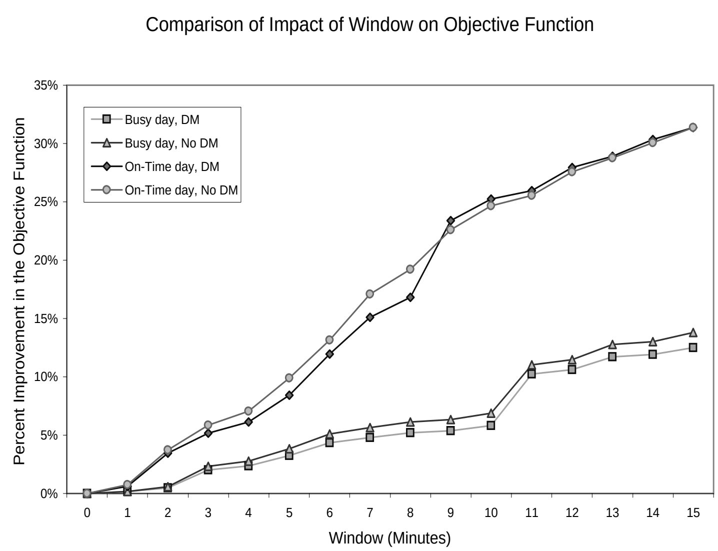


Figure 6-10: Impact of the inclusion of the delay multiplier on the benefits of preferential sequencing

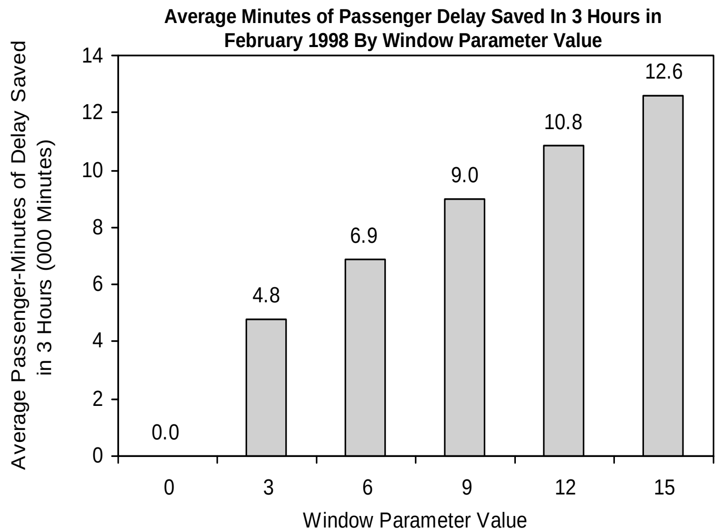


Figure 6-11: Average passenger minutes of delay saved during a 3-hour period in February 1998

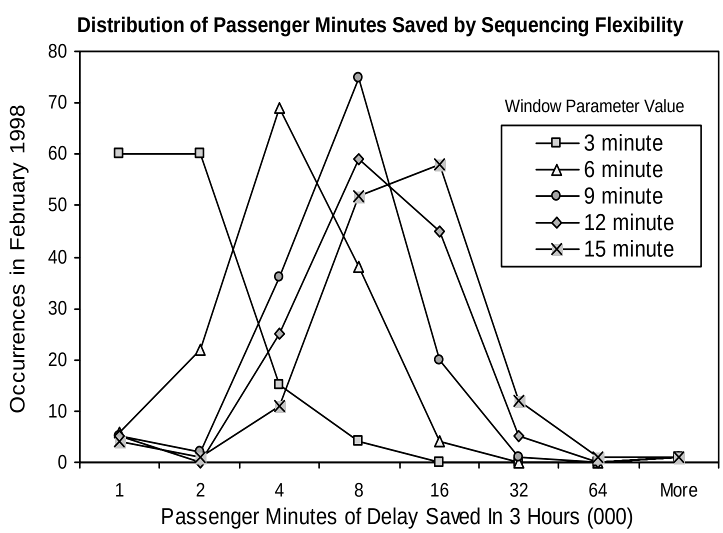


Figure 6-12: Distribution of passenger minutes of delay saved during a 3-hour period in February 1998

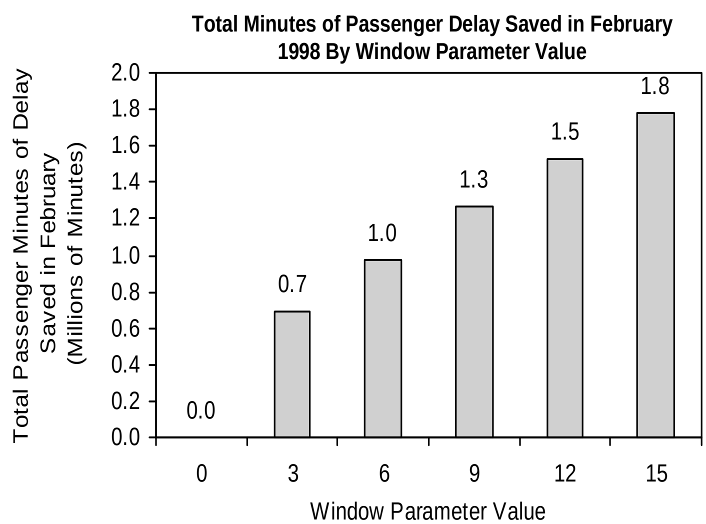


Figure 6-13: Total passenger minutes of delay saved during a 3-hour period in February 1998

# Chapter 7

# Conclusions

This study explores the benefits of increasing communication and collaboration between airlines and air traffic controllers during the arrival process at hub airports. In particular, this study estimates operational improvements, as measured by passenger minutes of delay incurred, from using CAP to provide airlines with more accurate landing time estimates and from allowing airlines to influence the sequence of their incoming traffic.

To estimate these potential benefits, the Arrival Sequencing Model was developed to simulate airline decisions regarding ground operations. This document presents the design, formulation, calibration and validation of ASM. The model has shown to solve relatively large problems quickly; a problem considering 60-80 aircraft over a time horizon of 3-4 hours solves in a matter of seconds. The efficiency of the model implies that ASM could be adapted for real-time decision making.

The results indicate that the potential benefits from increased communication and collaboration during the arrival process could be significant. Decreasing the standard deviation of the landing time estimate error from 5 minutes to 3 minutes could have prevented 500 passenger minutes of delay during a 3.25 hour period operating more or less on schedule, or 2000 passenger-minutes of delay during a 3.25 hour period experiencing significant delays. For the same time periods, allowing an airline to shift an aircraft's landing time by up to 6 minutes could save 5500 passenger minutes of delay in the On-Time period and 9200 passenger minutes of delay in the Busy time period. Furthermore, six minutes of sequencing flexibility could save 1 million passenger minutes of delay over one month or 12 million passenger minutes of delay over an entire year.

The preferential sequencing results indicate that further investigation into the feasibility of the sequences generated by ASM is warranted. Recall that ASM only limits the magnitude of the sequence change of an aircraft, meaning the number of minutes an airline is allowed to move up or back an aircraft's landing time relative to its original landing time. It considers neither airspace geometry constraints nor aircraft maneuvering limitations, such as speed and acceleration, when designing a sequence. Nor does it allow an airline to use slots originally assigned to another airline, even if the other airline is no longer using them. Therefore, the cited passenger minutes of delay saved from preferential sequencing may differ somewhat from the true potential. To better understand the limits of these improvements, a feasibility study of the arrival sequences generated by ASM is needed. This study will require information about an aircraft's positioning in the airspace surrounding an airport. It will also require input from air traffic controllers in order to define feasible sequencing moves.

# Bibliography


[1] Air Transport Association. *Airline Handbook*. Posted at website http://www.air-transport.org/public/Handbook/.

[2] Kari Andersson, Francis Carr, Eric Feron, and William Hall. Analysis and modeling of ground operations at hub airports. In *Third USA/Europe Air Traffic Management Research and Development Seminar*. EuroControl, June 2000. Posted at website http://www.eurocontrol.fr.

[3] Roger Beatty, Rose Hsu, Lee Berry, and James Rome. Preliminary evaluation of flight delay propagation through an airline schedule. In *Second USA/Europe Air Traffic Management Research and Development Seminar*. EuroControl, December 1998. Posted at website http://www.eurocontrol.fr.

[4] Bertrand Delcaire and Eric Feron. Dealing with airport congestion: Development of tactical tools for the departure flows from a large airport. Technical report, International Center for Air Transportation, Massachusetts Institute of Technology, 1998.

[5] DFW Airport Operations Department. DFW Statistics. Posted at website http://www.dfwairport.com.

[6] Robert Fourer, David Gay, and Brian Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. Boyd and Fraser Publishing Company, 1993.

[7] William Hall. *Efficient Capacity Allocation in a Collaborative Air Transportation System*. PhD thesis, MIT, January 1999.

[8] ILOG CPLEX Division. *AMPL User Guide, Including CPLEX Directives*, 1999.

[9] ILOG CPLEX Division. *ILOG CPLEX 6.6*, January 2000.

[10] Office of Aviation Policy and Plans, FAA. *Documentation for the Consolidated Operations and Delay Analysis System*, September 1997.

[11] Nicolas Pujet, Bertrand Delcaire, and Eric Feron. Input-output modeling and control of the departure process of busy airports. *Air Traffic Control Quarterly*, 8(1):1–32, 2000.

[12] Richard Zelenka, Roger Beatty, and Shawn Engelland. Preliminary results of the impact of CTAS information on airline operational control. Technical report, American Institue of Aeronautics and Astronautics, Inc, 1998.

# Appendix A

# ASM Formulation

$$
\begin{aligned}
Land_{ac,t} - Land_{ac,t+1} &\leq 0 \;\; \forall\, ac \in ArrPlanes, \\
&\qquad \forall\, t \in Time:\; t \geq FirstLand_{ac} \\
Gate_{ac,t} - Gate_{ac,t+1} &\leq 0 \;\; \forall\, ac \in ArrPlanes, \\
&\qquad \forall\, t \in Time:\; t \geq FirstGate_{ac} \\
Dep_{ac,t} - Dep_{ac,t+1} &\leq 0 \;\; \forall\, ac \in DepPlanes, \\
&\qquad \forall\, t \in Time:\; t \geq FirstDep_{ac} \\
Land_{ac,LastLand_{ac}} &= 1 \quad \forall\, ac \in ArrPlanes \\
Gate_{ac,T} &= 1 \quad \forall\, ac \in ArrPlanes \\
Dep_{ac,T} &= 1 \quad \forall\, ac \in DepPlanes \\
\sum_{ac \in ArrPlanes_t} (Land_{ac,t} - Land_{ac,t-1}) &\leq NumLand_t \quad \forall\, t \in Time \\
Land_{ac,t-Taxi_{ac}} - Gate_{ac,t} &\geq 0 \quad \forall\, t \in Time:\; t < T, \\
&\qquad \forall\, ac \in NoDep:\; FirstGate_{ac} \leq t \\
Land_{ac,t-Taxi_{ac}} - Gate_{ac,t} &\geq 0 \quad \forall\, t \in Time, \\
&\qquad \forall\, ac \in Turn:\; FirstGate_{ac} \leq t \\
Gate_{ac,t-Turn_{ac}} - Dep_{ac,t} &\geq 0 \quad \forall\, t \in Time,\; \forall\, ac \in Turn: \\
&\qquad \max(FirstGate_{ac} + TURN_{ac}, FirstDep_{ac}) \leq t
\end{aligned}
$$

$$
\begin{aligned}
Delay_{ac} - \sum_{t=FirstLand_{ac}}^{T} (Land_{ac,t} - Gate_{ac,t}) &= -Taxi_{ac} \quad \forall\ ac \in Turn \\
TempDelay_{ac} - \sum_{t=FirstLand_{ac}}^{T} (Land_{ac,t} - Gate_{ac,t}) &= -Taxi_{ac} \quad \forall\ ac \in NoDep \\
Delay_{ac} - TempDelay_{ac} &\geq 0 \quad \forall\ ac \in NoDep \\
Cnx_{ac1,ac2} - Gate_{ac1,t} + Dep_{ac2,t+Connect} &\leq 1 \quad \forall\ t \in Time, \\
&\quad \forall\ ac1 \in ArrPlanes, \\
&\quad \forall\ ac2 \in DepPlanes : \\
&\qquad CnxCost_{ac1,ac2} > 0 \\
\sum_{i=1}^{t-1} NumRes_{ac,i} - Dep_{ac,t} \times ResNeeds_{ac} &\geq 0 \quad \forall\ t \in Time, \\
&\quad \forall\ ac \in Planes : \\
&\qquad FirstGate_{ac} \leq t \\
NumRes_{ac,t} - ResMax_{ac} \times Gate_{ac,t} &\leq 0 \quad \forall\ t \in Time, \\
&\quad \forall\ ac \in Planes : \\
&\qquad FirstGate_{ac} \leq t \\
\sum_{ac \in DepPlanes_t} NumRes_{ac,t} &\leq Res\ \forall\ t \in Time \\
\sum_{ac \in DepPlanes_t} (Gate_{ac,t} - Dep_{ac,t}) \times SizeAC_{s,ac} &\leq NumGates_s\ \forall\ s \in Size, \\
&\quad \forall\ t \in Time \\
\sum_{ac:t \geq FirstDep_{ac}} (Gate_{ac,t} - Dep_{ac,t}) &\leq TotGates\ \forall\ t \in Time
\end{aligned}
$$

Note that in the formulation above, we are making the following assumptions:

1. $ArrPlanes_t$ is the set of all arrival aircraft with $FirstLand_{ac} \geq t$
2. $DepPlanes_t$ is the set of all departure aircraft with $FirstDep_{ac} \geq t$

3. $Land_{ac,FirstLand_{ac}-i} = 0 \quad \forall\ i > 0$

4. $Gate_{ac,FirstGate_{ac}-i} = 0 \quad \forall\ i > 0$

5. $Dep_{ac,FirstDep_{ac}-i} = 0 \quad \forall\ i > 0$

6. $Land_{ac,T+i} = 1 \quad \forall\ i > 0$

7. $Gate_{ac,T+i} = 1 \quad \forall\ i > 0$

8. $Dep_{ac,T+i} = 1 \quad \forall\ i > 0$

9. $\forall\ ac \in NoLand, Gate_{ac,t} = 1$ iff $t \geq GivenGateTime_{ac}$, which is the first time period the aircraft is at the gate

# Appendix B

# Feasibility Model

As discussed in Section 5.3, a second model was built to test the feasibility of the aggregate gate and resource assignment constraints in ASM. The feasibility model uses as input the aircraft movement times generated in ASM, namely the arrival at gate and pushback from gate times, to generate feasible gate assignments and ground crew team assignments. A feasible gate assignment means the aircraft is assigned a gate designed to service its aircraft type. Further, it means the aircraft remains at the same gate during its turn process. A feasible ground crew team assignment means a team of baggage handlers is assigned to the aircraft while it is parked at the gate. The crew assigned to the aircraft must also be part of the ground crew allotted to the zone in which the aircraft is parked. The formulation of the feasibility model is described here.

This formulation includes three sets of variables. The first, $Surplus_{z,t}$ indicates the number of extra resources needed to turn aircraft in Zone $z$ at time $t$. The second, $Assign_{ac,g}$ represents the gate assignment of the aircraft as follows:

$$Assign_{ac,g} = \begin{cases} 1 & \text{if } ac \text{ assigned to gate } g \\ 0 & \text{otherwise} \end{cases}$$

The third, $ResOn_{ac,t}$ indicates the number of resources assigned to aircraft $ac$ at time $t$. The parameters will be defined as they are introduced.

As discussed in Section 5.3, the objective function, defined in (B.1), is designed to minimize the number of surplus resources required to turn the aircraft. If an optimal solution exists with objective value 0, then we know that the movement times generated by ASM have feasible gate and resource assignments. If, however, an optimal solution exists with positive objective value, then we know that no feasible resource assignment exists. In particular, we know when and how many extra resources would be needed to get a feasible ground crew resource assignment.

The first three constraints consider physical gate constraints. Constraint (B.2) ensures that aircraft are assigned to compatible gates. Note that $Fit_{ac,g}$ is an indicator parameter having value 1 if aircraft $ac$ can fit at gate $g$, 0 otherwise. Constraint (B.3) ensures that no more than one aircraft is assigned to a particular gate at any time $t$. Note that $ACParked_t$ is the set of all aircraft parked at the gates at time $t$. Constraint (B.4) ensures that gate adjacency constraints are met. We define $g1$ and $g2$ as adjacent gates that have restrictions on the types of aircraft that can be simultaneously parked at the gates. We denote the set of aircraft parked at $g1$ that cause a restriction at $g2$ to be $Adj_{g1}$. The complete set of adjacent gate pairs with restrictions is denoted ADJ. The corresponding set of aircraft at $g2$ is denoted $Adj_{g2}$. For example, consider the previous example where gate 12A can service all types of aircraft, but when servicing a widebody, adjacent gate 12 becomes unusable. In this example, $g1$=12A and $g2$=12. Further $Adj_{g1}$ contains all widebody aircraft and $Adj_{g2}$ contains all aircraft.

The last four constraints consider the ground crew resource assignment problem. Constraint (B.5) ensures that the total resource-minutes assigned to the aircraft while it is parked at the gate meet or exceed the resources required to turn the aircraft. Note that $Turn_{ac}$ represents the minutes the aircraft is parked at the gate. Further, $ResNeeds_{ac}$ represents the resources required to turn the aircraft. Constraints (B.6)–(B.7) ensure that the appropriate resources are assigned, meaning resources assigned to the zone in which the aircraft is parked. Additionally, these constraints ensure that the number of resources assigned is within the appropriate range. Note that $Zone_g$ indicates the zone in which gate $g$ is included. Further, $MaxRes_{ac}$ and $MinRes_{ac}$ define the upper and lower bounds, respectively, on the number of resources that can be assigned to aircraft $ac$. Finally constraint (B.8) ensures that no more resources are assigned to gates in a zone than there are resources available in that zone. The number of resources available in zone $z$ is given by $Res_z$.

*MINIMIZE:*

$$\sum_{z \in ZONES} \sum_{t \in TIME} Surplus_{z,t} \tag{B.1}$$

*SUBJECT TO:*

$$\sum_{g \in GATES} Assign_{ac,g} \times Fit_{ac,g} = 1 \tag{B.2}$$

$$\forall\ ac \in \text{PLANES}$$

$$\sum_{ac \in ACParked_t} Assign_{ac,g} \leq 1 \tag{B.3}$$

$$\forall\ t \in \text{TIME}$$

$$\forall\ g \in \text{GATES}$$

$$\sum_{ac \in Adj_{g1} \cap ACParked_t} Assign_{ac,g1} + \sum_{ac \in Adj_{g2} \cap ACParked_t} Assign_{ac,g2} \leq 1 \tag{B.4}$$

$$\forall\ t \in \text{TIME}$$

$$\forall\ t \in \text{TIME}$$

$$\forall\ (g1, g2) \in ADJ$$

$$\sum_{z \in ZONES} ResOn_{ac,z} \times Turn_{ac} \geq ResNeeds_{ac} \tag{B.5}$$

$$\forall\ ac \in \text{PLANES}$$

$$ResOn_{ac,z} - \sum_{g \in GATES:Zone_g = z} Assign_{ac,g} \times MaxRes_{ac} \leq 0 \tag{B.6}$$

$$\forall\ ac \in \text{PLANES}$$

$$\forall\ z \in \text{ZONES}$$

$$ResOn_{ac,z} - \sum_{ginGATES:Zone_g = z} Assign_{ac,g} \times MinRes_{ac} \geq 0 \tag{B.7}$$

$$\forall\ ac \in \text{PLANES}$$

$$\forall\ z \in \text{ZONES}$$

$$\sum_{acinACParked_t} ResOn_{ac,z} - Surplus_{z,t} \leq Res_z \tag{B.8}$$

$$\forall\ t \in \text{TIME}$$

$$\forall\ z \in \text{ZONES}$$