Applications of ASDE-X Data to the Analysis of Airport Surface Operations

by

Melanie Sandberg

Submitted to the Department of Civil and Environmental Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Transportation

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012

Author ....
Department of Civil and Environmental Engineering
May 25, 2012

Certified by ....
Hamsa Balakrishnan
Assistant Professor of Aeronautics and Astronautics and Engineering Systems
Thesis Supervisor

Accepted by ....
Heidi M. Nepf
Chair, Departmental Committee for Graduate Students

Applications of ASDE-X Data to the Analysis of Airport Surface Operations

by

Melanie Sandberg

Submitted to the Department of Civil and Environmental Engineering
on May 25, 2012, in partial fulfillment of the
requirements for the degree of
Master of Science in Transportation

Abstract

While much attention has been given to analyzing and optimizing problems in air transportation, relatively little research has gone into studying airport surface operations. In recent years a surface surveillance system called Airport Surface Detection Equipment, Model-X (ASDE-X) has been installed at over 30 airports in the US as a safety device. The applications of the data being captured by these systems are far broader than just promoting safety. In this thesis, it will be demonstrated how ASDE-X data can be analyzed to characterize airport operations, and how it might be used going forward in real-time. The process of converting the raw ASDE-X data into a useable format will be discussed. Then, an analysis of airport operations at LaGuardia Airport and Philadelphia Airport will be presented using three months of summer data. These airports will be studied both in an aggregate fashion as well as for individual runway configurations. Finally, a case study of an Android tablet application will be presented as a next step in automation for aiding airport traffic operations.

Thesis Supervisor: Hamsa Balakrishnan
Title: Assistant Professor of Aeronautics and Astronautics and Engineering Systems

Acknowledgments

First and foremost, I would like to thank my advisor Prof. Hamsa Balakrishnan for all her help. Her guidance has been invaluable in shaping the direction of this research and refining it. Prof. Balakrishnan has offered her time, motivation and ideas to developing both my research skills and this project in particular. I am very grateful for the opportunity to be part of her team. I would also like to thank Prof. John Hansman and Dr. Tom Reynolds of MIT Lincoln Laboratory for their ideas, suggestions and feedback. Thanks also to Dr. Tom Reynolds for his help in providing the ASDE-X data.

Next, I would like to thank my colleagues who have been so dedicated to mentoring and collaborating with me. I am very grateful to Ioannis Simaiakis for suggesting research areas, sharing his knowledge of the air transportation industry and motivating me. His ideas and feedback were instrumental in developing the tablet application. Thanks also go to Harshad Khadilkar for his advice and for his assistance from his previous work with ASDE-X data. Vivek Panyam deserves a special thank you for all his help in developing the Android application. Many thanks also go out to all my labmates in the International Center for Air Transportation (ICAT) lab for their suggestions and support.

Lastly, I would like to thank my parents, family and friends for all their encouragement and support over the past two years.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Air transportation has long been a topic of interest to the research world, but only recently have airport surface operations been a subject of this research. The study of airport operations is of interest to airlines, airports and governmental authorities alike. More efficient operations on the surface can lead to both environmental benefits from minimized emissions and cost savings from reduced fuel burn. However, detailed data on aircraft movements on the surface was lacking until relatively recently. Across the US, over 30 airports have installed the Airport Surface Detection Equipment, Model X (ASDE-X) surveillance systems as a means to prevent runway and taxiway incursions. These systems track all aircraft on the airport surface, thereby providing position information not previously available. The use of this data to analyze airport operations is the main focus of this thesis.

1.2 Thesis Development

While the primary use of the ASDE-X surveillance systems is as a safety device, the vast amount of data collected enables new insights into airport surface operations. The precise one-second updates of aircraft locations can be used to do a historical analysis of airport surface movements, both at an aggregate level and specific to particular runway configurations, in order to characterize typical airport operations. The XML format for delivery of the ASDE-X data entries was optimized to reduce the amount of bandwidth needed to transmit the data. As will be described in Chapter 2, this is useful for data transfer, but requires extra effort to fill in all the missing pieces before it is usable for processing and analysis.

Chapter 2 details the format and contents of ASDE-X data in its raw form. It describes a method for sorting flights by tracks and then populating the empty data fields. After further pre-processing, the latitude and longitude coordinates for these sorted flights are used to create airport surface visualizations. These visualizations are then used to define queue boxes and to observe taxi operations based on runway configuration. Further processing yields a set of relevant metrics according to flight and also according to time period.

In Chapter 3, the daily outputs from the ASDE-X data processing are analyzed over a period of several months to describe typical operations at LaGuardia Airport. This includes metrics like departure and arrival counts, and average taxi times. Airport visualizations are used to show runway queuing examples and define queue boxes.

This study is continued by breaking down the analysis based on runway configuration. In Chapter 4, the four most frequently used runway configurations at LaGuardia Airport (LGA) are presented. For each configuration, analyses such as average taxi times and departure rates are presented. To conclude the chapter, the differences between the configurations are discussed.

To illustrate some of the differences between airports, a study of Philadelphia International Airport (PHL) is presented in Chapter 5. This aggregate analysis covers the same topics as before and highlights the differences between PHL and LGA. Chapter 6 presents the same metrics considered in Chapter 4, but for the two main runway configurations at PHL. The two runway configurations are contrasted with an analysis of the differences, and are then compared with LaGuardia Airport.

Chapter 7 discusses some next steps towards future uses of ASDE-X data. A case study is presented of a tablet application used in a trial at Boston Logan Airport. It demonstrates one step in the potential to use ASDE-X data to build decision-support tools for air traffic controllers. Finally, Chapter 8 summarizes the important findings of this thesis.

Chapter 2

ASDE-X Data

2.1 Introduction

2.1.1 Overview

The Airport Surface Detection Equipment, Model X (ASDE-X) system is a safety device that has been installed at over 30 airports nationwide as a means of preventing runway and taxiway incursions on the airport surface. The ASDE-X equipment gathers position data on all aircraft operating on the surface and uses it to alert air traffic controllers of potential conflicts and safety concerns. The position information is an agglomeration of surface radar and aircraft transponders. The ASDE-X system correlates the positional and identification data from the multiple sensors available at the airport and provides one second position reports on each track of data. A track contains all the fused data for each vehicle and aircraft on the surface or within a certain distance of the airport [2]. Fields being collected include the track number, callsign, aircraft type, altitude, latitude and longitude. From these measurements, other metrics such as velocity, x position, y position, and acceleration are derived.

An XML dataset provided by the Volpe National Transportation Center is used in this thesis. The particular data that is utilized by this distribution system is the Category 11 System Track Reports (CAT-11). The FAA currently archives data for 15 different airports around the country all interspersed and sorted by timestamp.

2.1.2 Structure

There are two main types of messages in the XML data distribution format: system track reports and status reports. The status reports provide information about the components of the ASDE-X system such as mode and quality. The system track reports provide information about the aircraft and vehicles being tracked [4]. Within a given dataset, the majority of entries are system track reports for individual aircraft. An example of a single system track report from the raw XML file is shown in Figure 2-1. The message body contains the relevant position information needed for this analysis.


```
2011-06-01T00:00:42

BytesMessage[ID:<306650.1306886399981.0>]

Property: id
Value: KPHX.1466

Property: airport
Value: KPHX

Property: sendTo
Value: all

Property: msgType
Value: positionReport

Property: DEX_SOURCE_TYPE
Value: ASDEX

Property: airline
Value: DAL

Property: JMSXDeliveryCount
Value: 1

Message Body (size = 674): <asdexMsg xmlns="urn:us:gov:dot:faa:ato:asdex-sd">
<airport>KPHX</airport><positionReport><time>2011-05-31T00:00:00.000Z</time><track>
1466</track><position><x>463</x><y>-349</y><latitude>33.4321703</latitude>
<longitude>-112.00089924</longitude><altitude>1131</altitude></position><movement>
<speed>15</speed><heading>91.8</heading></movement><status><gbs>1</gbs><src>default
</src></status><targetExtent><startRange>1972</startRange><endRange>2017</endRange>
<startAzimuth>59.3</startAzimuth><endAzimuth>60.4</endAzimuth></targetExtent>
<plotNumber><plot1>34546814</plot1><plot2>1887510</plot2><plot3 rdr="smr1"
src="mlat">1887524</plot3></plotNumber></positionReport></asdexMsg>
```


Figure 2-1: Sample system track report

The reports are listed by time, with new measurements being recorded every second, causing all the aircraft tracks to be intermingled with each other. The ASDE-X CAT-11 binary data generally contains a complete data set for each target at every second, potentially sending the same values multiple times. The high bandwidth needs of the XML distribution makes it difficult to transfer this volume of data, even after compression. Therefore, in order to limit the amount of data being transferred, not all fields within the position report messages are fully populated.

Each aircraft XML position report is either a full report or a partial report. A full message is sent on the first occurrence of a track and at a configurable frequency, for example every sixty seconds, after that initial message. Partial messages are sent in the updates between the full ones, and contain identification information and only the fields that have changed since the previous message [4]. This approach significantly reduces the bandwidth required for the XML feed. Some data fields appear in every position message: Examples include the airport that generated the message, the track number at the airport, and the time of the message.

2.2 Data Preparation

The pre-processing phase of the ASDE-X analysis using Matlab includes loading in the data, filling in missing values, sorting flights by callsign, defining flights as departures or arrivals, dividing callsigns into multiple flight legs, removing en route flights and then performing filtering on the data.

2.2.1 Extraction of Relevant Fields

The XML aircraft system track reports in their raw form contain some redundant information. Preparing the raw data for analysis requires filtering this data, first by relevant message information, and then by airport. In this case what remains is the position information from all the system track reports for aircraft at a given airport.

Because the process of extracting the compressed XML files is time-consuming and because the extracted files are quite large, there is a need for a system that automatically parses the data and returns only the relevant information. It is also foreseeable that we may want to analyze other airports in the future. Therefore, an initial file containing information for all airports is created, which could be used later for studying any of the other 15 airports contained in the data. This makes the process more efficient by removing the need to repeat the time-consuming data extraction process. A script was written which reads in one of these XML files, parses through the message bodies, and writes only the relevant fields to a new file. In our analysis the relevant fields are: Airport, timestamp, track number, longitude, latitude, speed, heading, callsign and aircraft type.

This file, which contained information from all airports, was archived for possible future use. A second script was written to read in this archived file, search for messages from the chosen airport, and then write these entries to a new file. These files are considerably smaller which speeds up the processing time. By reducing the data down to just one airport, the Matlab code does not need to read in and parse irrelevant data. The final product is a .csv file similar to that shown in Figure 2-2.

	A	B	C	D	E	F	G	H	I	J
1	Airport	Timestamp	TrackNum	Lon	Lat	Speed	Heading	Callsign	Aircraft Type	
80	KBOS	2011-06-01T00:00:37.000Z	2897	-71.0184	42.3658					
81	KBOS	2011-06-01T00:00:37.000Z	3009	-71.0181	42.3655					
82	KBOS	2011-06-01T00:00:37.000Z	2936	-71.0126	42.3577	6	206.4	AAL108	B752	
83	KBOS	2011-06-01T00:00:37.000Z	2493	-70.9374	42.3691	185	80.1			
84	KBOS	2011-06-01T00:00:37.000Z	3092	-71.0144	42.3639					
85	KBOS	2011-06-01T00:00:37.000Z	2765	-71.0072	42.377		197.5			
86	KBOS	2011-06-01T00:00:37.000Z	2715	-71.0045	42.3744		200.8			
87	KBOS	2011-06-01T00:00:37.000Z	2976	-71.0144	42.361		197.9			
88	KBOS	2011-06-01T00:00:37.000Z	3055	-71.0137	42.3473					
89	KBOS	2011-06-01T00:00:37.000Z	3017	-71.0126	42.3558	3				
90	KBOS	2011-06-01T00:00:37.000Z	3082	-71.0173	42.355		195.7			
91	KBOS	2011-06-01T00:00:37.000Z	2653	-71.0184	42.3658					
92	KBOS	2011-06-01T00:00:37.000Z	3393	-71.0227						
93	KBOS	2011-06-01T00:00:37.000Z	3069	-71.0147	42.3587	6	109.6			
94	KBOS	2011-06-01T00:00:37.000Z	3111	-71.0186	42.3586					
95	KBOS	2011-06-01T00:00:37.000Z	2982	-70.704	42.2614	390	182.8			
96	KBOS	2011-06-01T00:00:37.000Z	2954	-71.007	42.361	145	19.9			
97	KBOS	2011-06-01T00:00:38.000Z	2493	-70.9363	42.3693	182	80.2			
98	KBOS	2011-06-01T00:00:38.000Z	2765	-71.0072	42.3769		194.8			
99	KBOS	2011-06-01T00:00:38.000Z	2839	-71.0181	42.3655	0				
100	KBOS	2011-06-01T00:00:38.000Z	2897	-71.0184	42.3658					
101	KBOS	2011-06-01T00:00:38.000Z	3053	-71.0226	42.3602	11	139.3			
102	KBOS	2011-06-01T00:00:38.000Z	2976	-71.0144	42.361	9	195.7			

Figure 2-2: Sample ASDE-X data file for LGA

2.2.2 Sorting

Because of the nature of the XML data, two main things are needed to make it usable for analysis. First, the data needs to be sorted and separated according to individual flights. Since not every system track report includes a reported value for the callsign, the flights are sorted by track number. Track numbers are reused over the course of the day so that alone cannot be the sole identifying factor for identifying the entries for a single flight. After sorting by track number, each track number is divided into multiple flights, if necessary, based on the time, callsign and location.

The easiest way to separate a track into multiple flights is by detecting a gap in time stamps between two successive entries. There are cases in which a flight track might not be detected for some period of time, because of equipment malfunctions or other glitches in the system, so a gap in time between successive entries does not necessarily imply the next entry is a different flight. Based on manually observing the data, it was observed that a typical gap in time between the overlapping use of track numbers was over 6 hours. It was also observed that most short gaps in the flights tracks were less than 10 minutes in length. Therefore it was decided that any gaps in the flight track of less than 10 minutes do not imply a change in flights. In addition to the time duration, the other relevant data field is the callsign. If the callsign for a given track number changes at any point, the first appearance of the new callsign is considered to be the start of a new flight record. After sorting and filling (described in the next section), all tracks with an unknown callsign ("UNKN") are discarded. Many of these tracks correspond to ground vehicles.

2.2.3 Completing Missing Fields

After sorting by track number, the other measure needed to produce meaningful data is to fill in all the fields which have been left blank. As stated above, many fields are only populated when there has been a change in its value in order to prevent sending redundant information. After flights have been grouped together by track number, each flight's data is sorted by timestamp and the entry is processed. As each system track report is read, if a field is missing from that message, the corresponding value from the previous message is used to populate the field. This is only done when the previous and current message have the same track number and when the time interval between the two messages is less than 10 minutes.

Finally, after all the empty fields have been populated, the issue of short gaps in time still exists. Those gaps that are less than 10 minutes are not considered to be a change in flight legs, yet no data is available about the aircraft in that time. In the absence of data, the aircraft will not be counted in the sum of total aircraft on the ground, or in the length of a runway queue. For the sake of determining metrics like taxi time and time in queue, message entries are generated for each second of missing data. Because there is no simple way to extrapolate surface information based on the position before and after the gap, the last known position, along with all other last known data fields, was replicated for each of the generated timestamps. While it might be argued that it would be better to not generate inaccurate data for the missing time period, it is necessary to have some record of the flight during that time period, even if it inaccurate. For purposes of determining how many aircraft are active on the surface at any given point in time, it is useful to have a placeholder aircraft, even if its exact position may not be known.

2.2.4 Classification of Flights

Flights are defined as departures or arrivals based on their initial and final locations. If the last entry is further away from the control tower than the first entry, then it is classified as a departure. If it is closer, then it is classified as an arrival. While all flights are classified as one or the other, there are a few exceptions, such as overflights, surface movements and multiple flight legs.

1. **Overflights:** Some of the flights registered are actually overflights or flights to neighboring airports, which never actually depart or land from the chosen airport. These will be discussed below.

2. **Surface Movements:** Some flights are on the airport surface but never depart. These flights taxi out and then at some point turn around and return to the terminal. This could be because of a mechanical problem or because of weather or because of the 3-hour tarmac rule.

3. **Multiple flight legs:** There will be a brief discussion below of how flights of the same callsign are divided into multiple flights. However, this does not catch every occurrence of the same aircraft both arriving and then taking off because there must be a gap of at least 10 minutes between successive entries. If the ASDE-X data keeps

recording entries for the aircraft every minute it will not be divided into multiple flight legs and either the arrival or the departure leg will be missed. This was seen in some of the general aviation aircraft and may be because of the transponder not being turned off. In this case, the way to try to determine when the arriving flight leg ends was to look for a flight that has stopped for more than 20 minutes.

4. **Separating Flight Legs:** The full set of entries for a single callsign may actually be composed of several flight legs. If a gap in time between two successive entries is greater than 10 minutes, these will be classified as separate flight legs.

5. **Removing Overflights:** For the purposes of this analysis, only flights that are at some point active on the surface of the chosen airport are relevant. Therefore any overflights are removed from the stored dataset. A flight is considered an overflight if its velocity is never less than 25 knots.

The final step in the data preparation is Unscented Kalman Filtering, which is performed to remove the noise from the aircraft surface tracks. It is discussed at length in Chapter 2 of Khadilkar's thesis [6].

This final product of fully-populated data fields, is used as the starting point for further data processing and analysis. The Matlab code being used for this analysis was originally created by Harshad Khadilkar for work on his master's thesis [6]. The code has been extended for the purposes of analysis for this work. The results of such data processing are described in Chapter 3 for LGA, and in Chapter 5 for PHL.

2.3 Airport Surface Visualizations

One of the great benefits of ASDE-X data over other data sources, such as ASPM data, is its ability to track aircraft locations throughout the whole of the taxiing operation. Using just the timestamp, latitude and longitude, it is possible to create visualizations of aircraft movements at the airport. For this study, the professional version of Google Earth™ was used to display the visualizations. Google Earth™ reads in .kml files which provide data entries of flight timestamps and geographical coordinates.

By incorporating some of the other ASDE-X data like callsign and aircraft type, it is possible to plot just a single aircraft, or to highlight a single aircraft amongst all the aircraft on the surface, perhaps by making it a different color or size. In the same way it is possible to change the color or size of all aircraft of a certain type. Figure 2-3 below shows an example of a Google Earth^{TM} screen-shot for LaGuardia Airport. In this case, departures are shown in green, and arrivals in red. These screen-shots can be animated over a period of time to see the progression of aircraft movements on the surface.

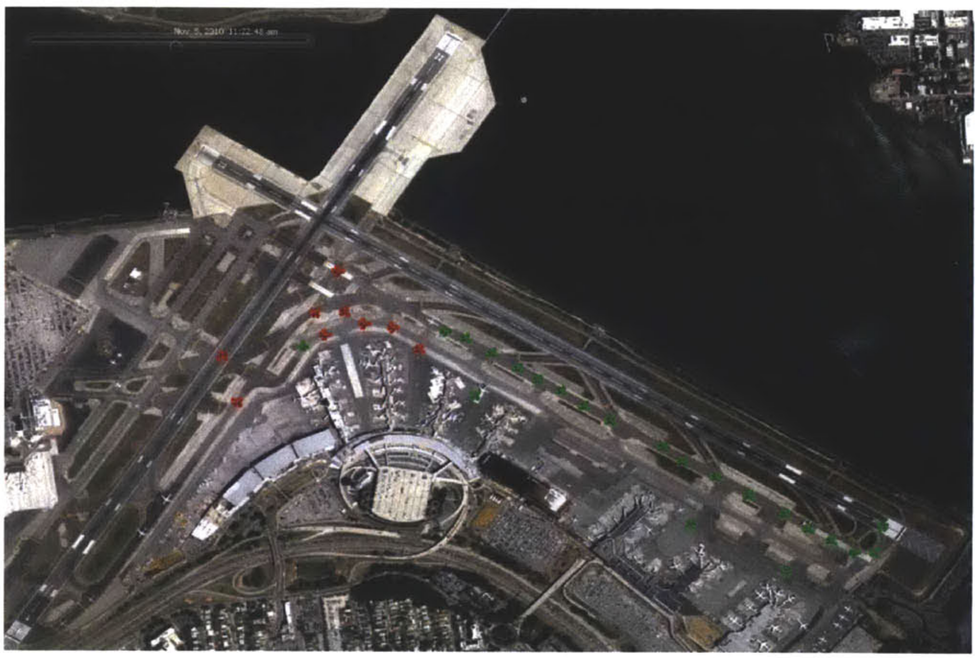

Figure 2-3: Runway threshold queue at LGA

Visualizations of airport surface operations have many uses. First, they are useful for determining queue boxes for the departure runways and also for runway crossings. Next, they can be used for determining where aircraft with extensive ground holds or EDCTs are asked to wait. Finally, they can be used to determine the taxi routes of departures. These features will be discussed in general terms here, and more specifically for LGA in the next chapter.

2.3.1 Queue Box Definitions

Queue boxes are used to determine the length of queues waiting for the departure runway, and to determine the amount of time that an aircraft spends waiting in line for the runway. At some airports, these are fairly intuitive to define. Some departure runways only have one taxiway leading up to it. Other runways are more complicated: They may have multiple taxiways leading to the runway threshold that are all used as part of the queue. Sometimes, the location of the queue for a given runway is dependent on which arrival runway is being used. Figure 2-3 above shows a departure queue for Runway 31 at LGA, with departures(in green) lined up in a single straight line for the runway.

By looking at visualizations, runway crossing queue boxes can also be determined. The length of these queues can be observed and taken into account in the data analysis. For example, if the number of aircraft waiting to cross the runway is always just one or two aircraft, then it is probably not part of the departure runway queue. The aircraft are only stopped because they are waiting for an arrival to pass, and can then proceed on with their normal taxi route. In this case, the time spent waiting to cross the runway should not be included in the count of the time spent in the departure queue. In other cases, the queue waiting to cross an active runway is, in fact, a part of the runway threshold queue. The observation of visualizations is the quickest way to determine which is the case. Figure 2-4 shows a runway crossing queue with departures headed for Runway 13. Part of the queue is at the threshold for Runway 13, with the rest of it continuing on the other side of the active Runway 4/22.

A further use of airport visualizations is to determine where aircraft are commonly held for EDCTs and other ground holds. This can easily be observed in the visualizations. One danger of not explicitly noting this is the potential to include these flights as part of the queue for the runway. This would affect both the counts of aircraft in the queue, and the average amount of time spent in the queue.

Finally, these visualizations are helpful to observe the path of a taxiing aircraft, which will vary based on the runway configuration and the initial location. It is useful to see where aircraft stop along the way, and for what reason. This visualization also gives a

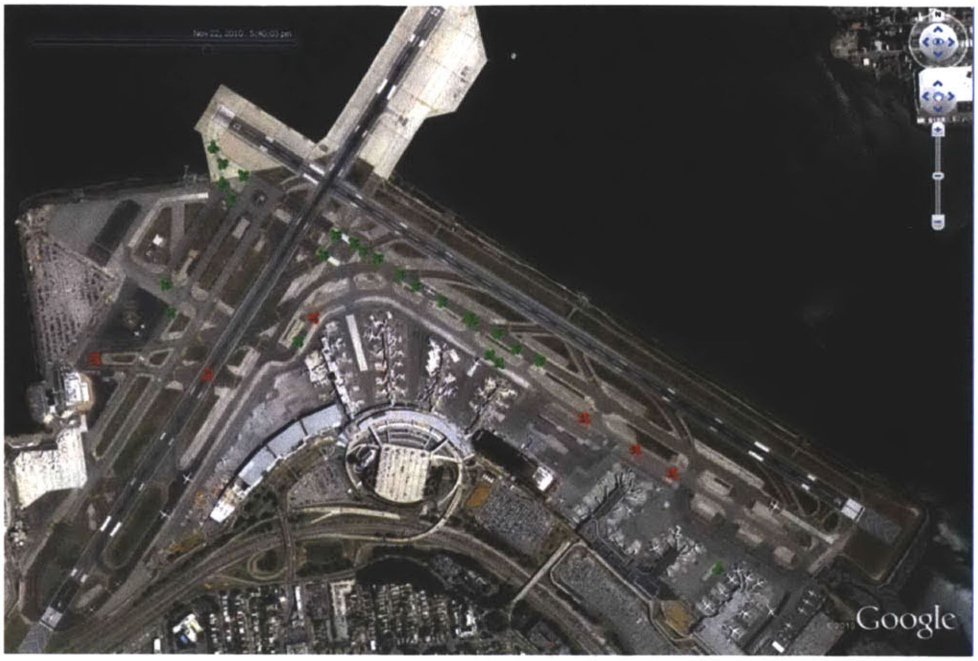

Figure 2-4: Runway crossing queue at LGA

good picture of the interaction between arrivals and departures. In particular, the number of arrivals and the primary arrival runway can significantly affect the departure taxi times, as a result of these interactions.

2.4 Data Processing

The rest of the Matlab code performs the main airport-specific processing. It takes the final data set produced by the pre-processing phase and analyzes it to provide meaningful outputs from the ASDE-X inputs. Different queue boxes are defined and aircraft are being tracked for the times that they enter and exit the queue boxes. Runway configurations are also defined in the main processing phase, and the code tracks which configurations are used over the course of a day. There are a wide variety of metrics that could be produced from this data. In the next section, some of the most common metrics that are being compiled will be defined.

2.4.1 Metrics

This section lists some common metrics used to characterize airport surface operations. For each metric, a description is included of how it has been defined for the purposes of studying LGA is presented to illustrate the concept.

1. **Taxi times:** The taxi time for a departure (i.e., the taxi-out time) is defined as the difference in time between initial detection and the time that the aircraft enters the runway. The initial detection is not necessarily the same as the pushback time, since it is likely that the aircraft transponder is not turned on until after pushback has already occurred. The resultant difference between the OUT time and the ASDE-X detection time will be studied in the next chapter, with a comparison of ASPM and ASDE-X push times. To account for aircraft that may have pushed back but are held on the ground for some extended period of time and are not actively taxiing, any aircraft that is stationary for more than 20 minutes is considered to be holding. The hold times are calculated based on the position of the aircraft not changing, or by a velocity less than 2 knots. While it might seem reasonable to require a velocity of zero knots, the noise in the data shows that a stopped aircraft may be reporting a velocity of up to 1.5 knots.

   The arrival taxi time is defined as the time difference between the wheels-on time of the aircraft and the final detection of the aircraft. Most aircraft exit the runway quickly after touchdown, so the use of the wheels-on time is deemed sufficient for a start to the taxi time. The arrival taxi times are typically much more consistent than the departure taxi times. However, given the level of congestion at LGA, the study of arrival taxi operations is an interesting examination.

2. **Queue Length:** Queue length is defined as the number of aircraft within the queue box for an active departure runway at a given time. As will be seen in the section on queue boxes, some runways have multiple queue boxes, and the resultant queue length will be the sum of aircraft in any queue boxes that are active.

3. **Time in Queue:** The time in queue is the amount of time that a given aircraft spends

within the queue box for the departure runway. This is measured by subtracting the time that the aircraft entered the queue box from the departure time.

4. **Departure Time:** The departure time is the closest approximation to the time that the aircraft lifted off the ground. It is measured by detecting the time that the aircraft velocity first exceeded 35 knots.

5. **Arrival Time:** The arrival time is the closest approximation to the time the aircraft touched down on the ground. This time is calculated by finding the earliest time that the aircraft velocity was less than 35 knots.

6. **Departure and Arrival Counts:** The number of departures and arrivals is tracked for each 15-minute interval throughout the day. Additionally, the time departure and arrivals times of each aircraft are tabulated for a more detailed description of the departure and arrival counts. These can also be used for determining inter-arrival or inter-departure times.

2.5 Outputs of Analysis

The main processing finishes by outputting several graphs which visually display many of the metrics defined above, over the course of the day. This helps researchers to visually observe peak demand times for departures and arrivals, and also peak taxi times. A quick analysis of several days can help determine whether these peaks are consistent over many days, or perhaps whether they vary by day of week. These graphs also make it easy to observe changes in runway configurations, and to track queue lengths over the course of the day. Figures 2-5 and 2-6 show some of the graphs produced from the ASDE-X data.

Figure 2-5 is divided into three parts. The top graph shows runway utilization. The overall height of the bar shows what percentage of the time the runway is being utilized during a 15-minute time period. The various colors within the bar show the breakdown of the ways the runway is being used in that time period. The middle graph shows the number of departures and arrivals for each 15-minute time period on the given runway. The bottom graph shows the length of runway departure queues. This will be empty if the given runway

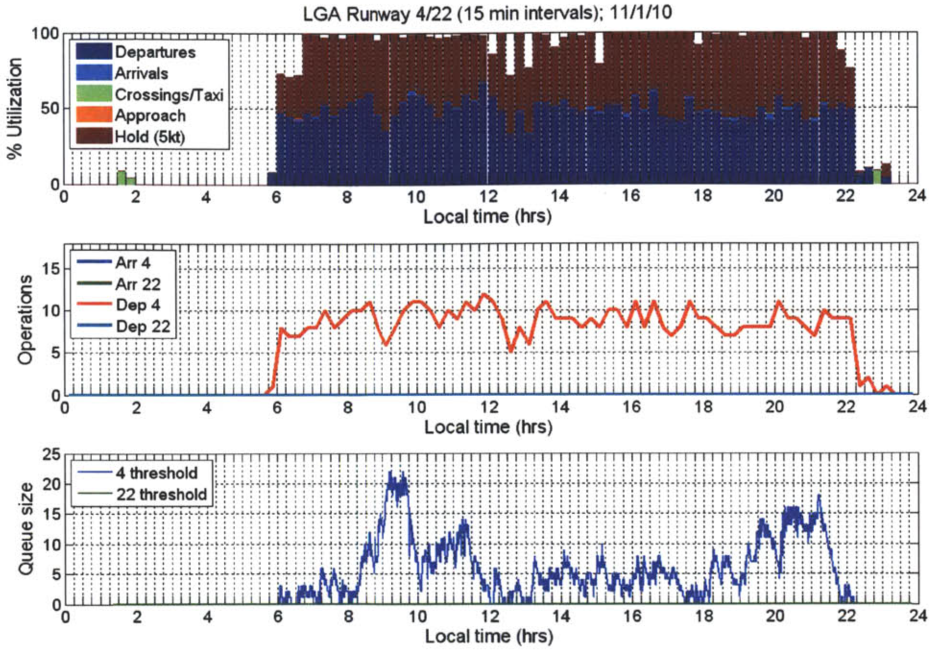

Figure 2-5: Runway 4/22 usage in LGA (11/1/2010)

is not being used for departures. These plots are created for each runway at a given airport. Figure 2-6 shows both the average taxi times and the number of aircraft that began taxiing in that 15-minute time interval.

The final result of the processing is a file which stores all the flight data, taxi times, runway configurations, departure and arrival counts and queue lengths. While these files alone contain volumes of useful data, when further analyzed over many weeks or months, an even better picture of airport operations can be developed. With that amount of data, the operations metrics can be analyzed both for the airport overall, and separately based on runway configuration.

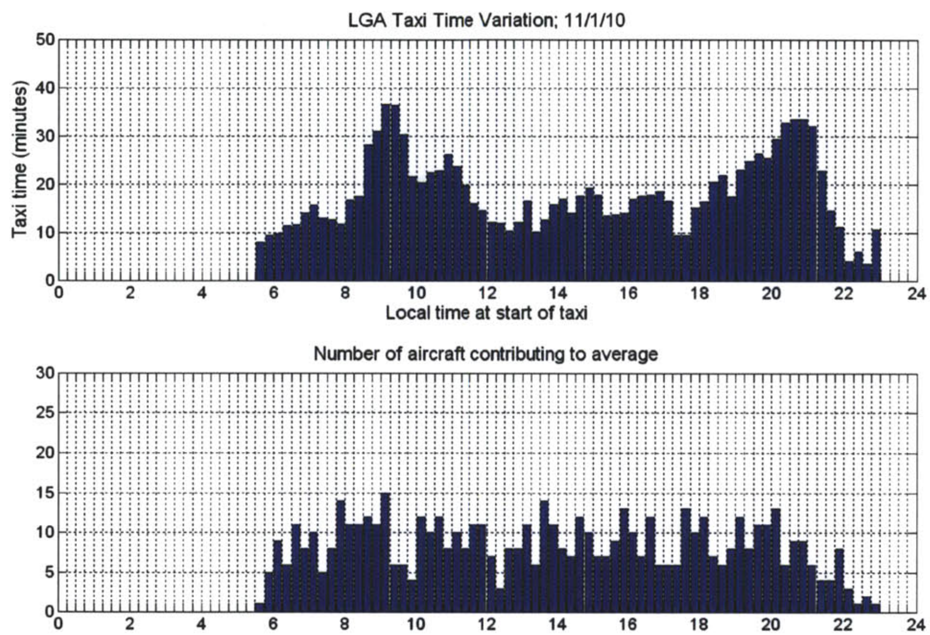

Figure 2-6: Average taxi-out times for LGA (11/1/2010)

Chapter 3

Analysis of Surface Operations at LGA

3.1 Introduction

LaGuardia Airport (LGA) is one of three major airports (along with JFK and EWR) in the New York metroplex area. It is the smallest of the three airports, with over 30,000 flights arriving and departing each month, and serving 24 million passengers per year. LaGuardia has four terminals: the Marine Air Terminal, the Central Terminal Building, the USAirways Terminal and the Delta Terminal. The Central Terminal Building has four concourses and over 40 gates. The Marine Air Terminal is for general aviation and Delta Shuttle operations. As can be seen in Figure 3-1, there are two intersecting runways that are in use, Runway 4/22 and Runway 13/31. In an effort to control congestion, it is one of the few airports in the US which has slot-restrictions on arrivals and departures.

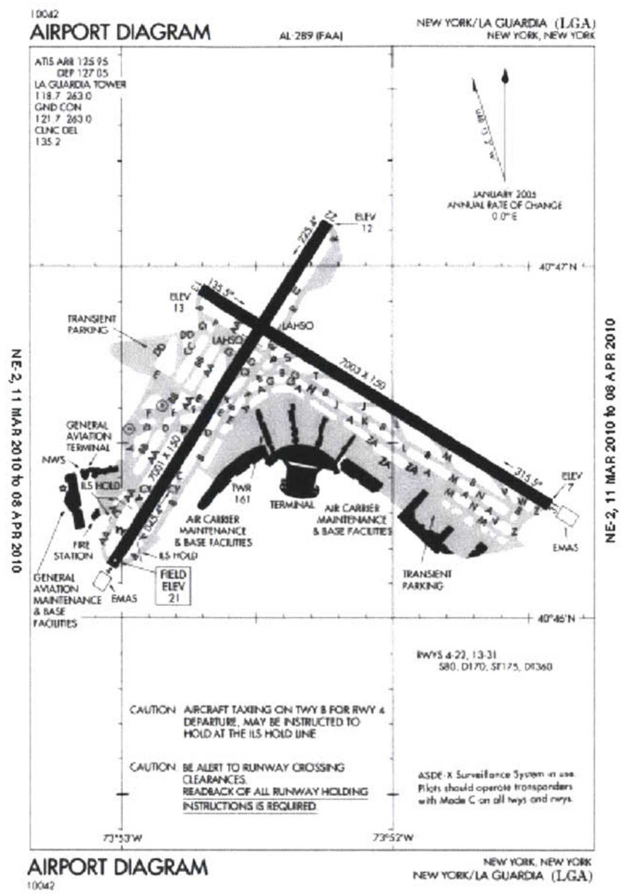

Figure 3-1: LaGuardia Airport Diagram

3.2 Motivation

The motivation for focusing on LaGuardia Airport is the magnitude of surface congestion that it experiences. An analysis of the ASDE-X data for LGA over a significant period of time is useful for characterizing normal operations, and also to evaluate whether surface flow traffic management initiatives might be helpful in reducing congestion. The ASDE-X data analysis and surface visualizations can be used to determine average demand, queue lengths, taxi time, and time in queue. The ASDE-X data can also be used to measure actual throughput and compare it to predicted capacity. Any significant differences between the predicted and actual values can be used for better airport and FAA planning. The data can also be used to determine how these metrics vary based on the configuration, time of day, or day of the week. While average values of these metrics are helpful in making predictions, it is equally useful to know the variances, and when the minimum and maximum values are typically achieved.

It is only by studying current conditions that baseline capacity values can be determined and then tracked over time. Specifically for the study of airport surface congestion, the taxi times and queue lengths are metrics that need to be monitored for any significant shifts.

In studying congestion, the goal is not just to quantify and predict metrics like taxi time and queue length, but to determine the causes behind them. The obvious cause is an excess of demand, but other more subtle causes of congestion may exist. For example, congestion in certain runway configurations may arise due to aircraft waiting to cross an active runway, or from an excess of arrivals. Once the causes have been determined, the best solution or system of solutions can be developed.

3.3 Comparison of Pushback Times

One of the metrics that is derived from the ASDE-X data analysis is taxi time, but this cannot be done without accurately knowing the pushback time. One problem with ASDE-X data is that the flight track only begins once the aircraft transponder has been turned on, which often occurs after the aircraft has pushed back from the gate. In this case, there is no way of knowing when the aircraft actually pushed back from the gate. Many previous studies of taxi times have been completed using Aviation System Performance Metrics (ASPM) data [1], which records a push time and a wheels-off time. The taxi time can be inferred as the difference between those two numbers. We conduct an analysis comparing the pushback or OUT time recorded in ASPM data and pushback or detection times reported in ASDE-X data. Once an average difference between the two data sources has been determined, the ASDE-X push times can be adjusted accordingly.

The pushback time in the ASDE-X data is considered to be the first recorded message for that particular aircraft track number. The push time in ASPM data is recorded under the field "ACTOUTTM". In ASPM data, two different types of pushback times may be recorded. One is an automatic recording of the pushback time, based on when the pilot releases the brakes of the aircraft. In this case another field, called "OOOI", is marked Y. OOOI stands for Out, Off, On, and In, and records the respective times for each one of those actions for a given flight. The other is a estimate of the pushback time and is marked N in the "OOOI" field. Analysis was done first using all the ASPM data and then using only the flights with automatic recordings (i.e., the OOOI flights), for the three months of summer 2011 in LGA.

The analysis was completed by comparing individual flight push times in ASPM and ASDE-X. For each day of ASDE-X data, a list of callsigns was compiled for departing aircraft. Those callsigns were then searched for in the ASPM data to find a match for the given day. Some callsigns needed to be altered slightly to make up for the differences between the two datasets. For example, flight callsigns starting with 'USA' in ASPM are listed as 'AWE' in ASDE-X data. There is a 95% match of flights between the two datasets. For each matching flight, a record was made of the two push times. After all the data was generated, an analysis was completed of the differences between the two sets of push times. The results were grouped and the differences shown in Figure 3-2. A negative number signifies that the ASPM data had an earlier pushback time than the ASDE-X detection. It is expected that the ASPM number would be earlier because the ASDE-X data does not pick up the flight track until the transponder is turned on, which is typically after pushback.

It was also recorded for each flight whether the ASPM time was automatically recorded (OOOI = Y) or not and he same analysis was performed only using these flights. Automatically recorded flights were roughly 57% of all flights at LGA in Summer 2011. Figure 3-3 shows the distribution of time differences for all automatically recorded flights (OOOI flights) grouped into bins of 5 minutes each.

The average difference in push times for all flights was -6.4 minutes, with over 50% of flights having a time difference between 0 and 10 minutes. The average difference in push times for the OOOI flights was -7.7 minutes. In both cases the ASPM time was, on average, earlier. We believe that the estimate from OOOI flights is more reliable. In the study of taxi out times below, the average difference in pushback times of OOOI flights could be added to the average taxi out times to give a more accurate prediction of taxi-out times.

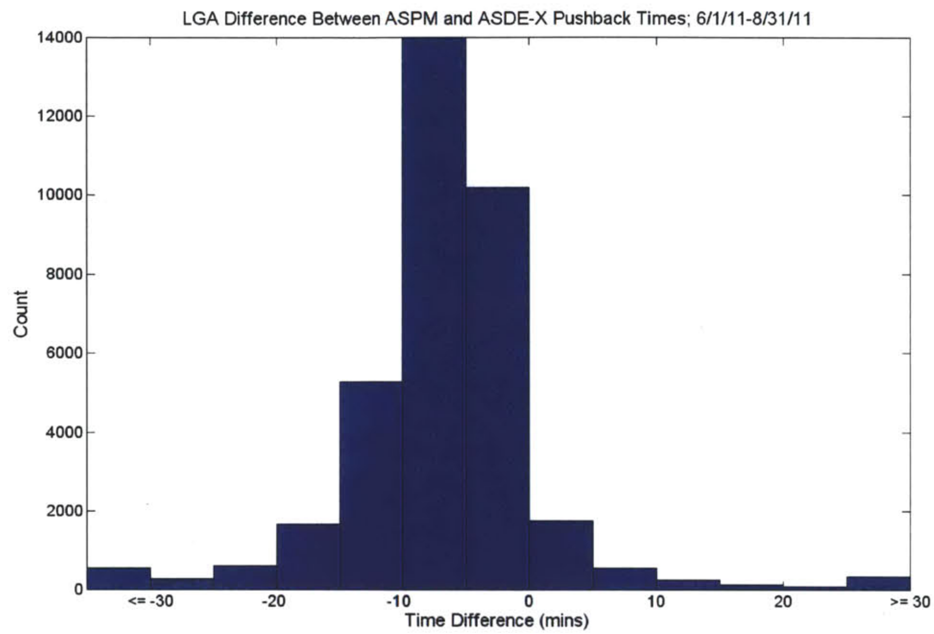

Figure 3-2: Difference between ASPM and ASDE-X push times; LGA Summer 2011

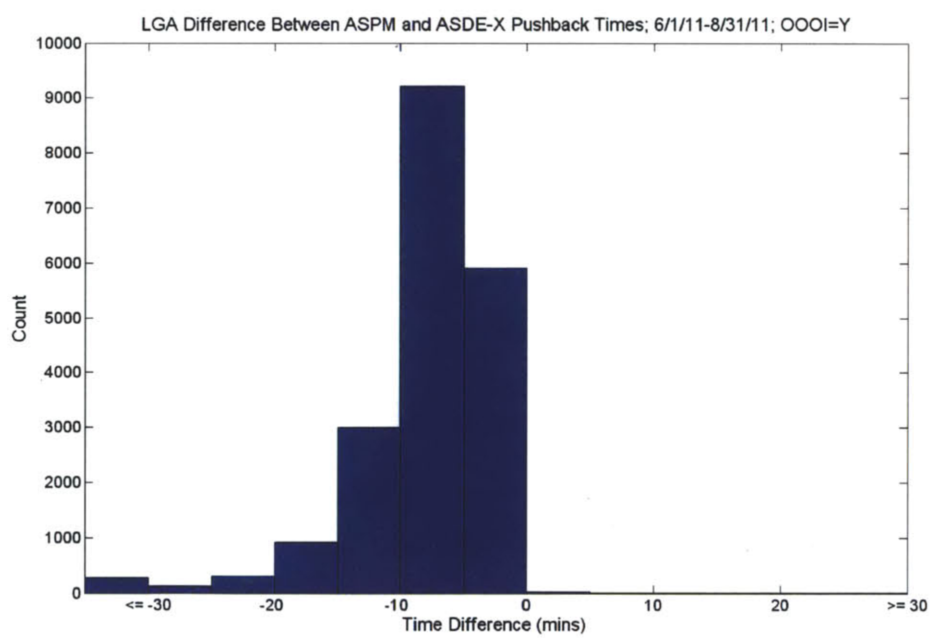

Figure 3-3: Difference between ASPM and ASDE-X push times when OOOI=Y (LGA Summer 2011)

3.4 Configuration Usage

LGA operates under four main runway configurations. This chapter attempts to characterize operations at LGA noting the differences between the various configurations. These are not the only configurations in use but these four were used the majority of the time in Summer 2011. Figure 3-4 below shows the percentages of time that each configuration was in use. The four most-used configurations in order are 22|13, 31|4, 22|31 and 4|13.

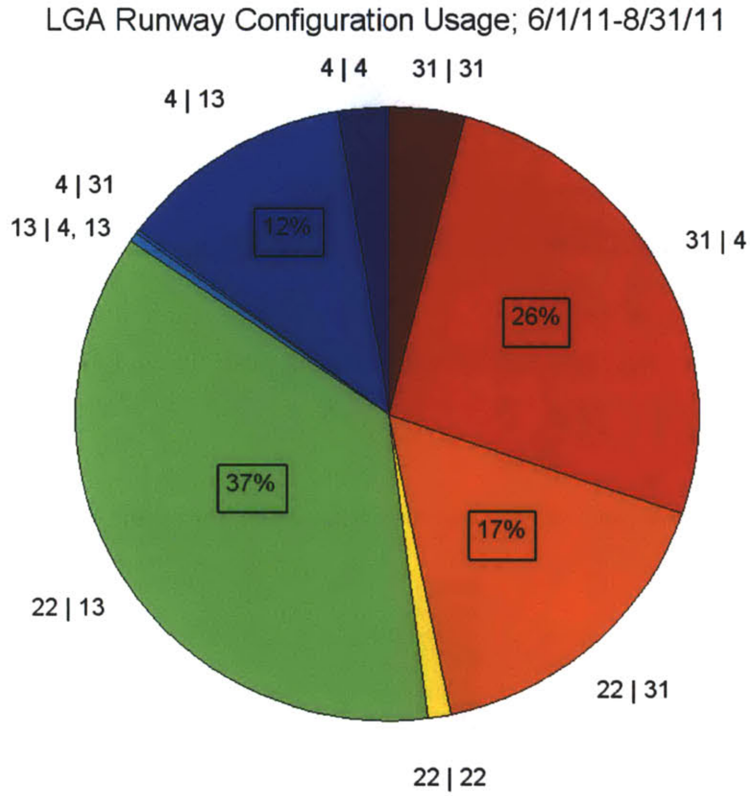

Figure 3-4: LGA runway configuration usage; Summer 2011

The configuration usage will vary by time of year, according to prevailing winds and demand patterns. This analysis only covers the summer months of 2011, and the configuration usage will likely differ in the winter months.

The most-used configuration was 22|13, which was used 37% of the time. Given LGA's two intersecting runways, it makes sense that the dominant configuration would have arrivals and from the end of the runway closer to the intersection. This way, the separation time between aircraft is minimized. Once an arrival has crossed the intersection, the next departing flight can begin its takeoff roll immediately. Similarly, there is less variance in the amount of time needed for a departure to clear the intersection, so the controller does not need to build in as much of a buffer for the arriving aircraft.

The remaining configurations are 31|4 (used 26% of the time), 22|31, (17% usage), and 4|13 (used 12% of the time). In these cases, the arrivals land at the far end of the runway, away from the intersection. When possible, arrivals use Land and Hold Short Operations (LAHSO) to minimize the wait time for departures, and also to avoid crossing an active runway during taxi-in.

In addition to the overall usage of each configuration, it is interesting to study the time of day that each configuration is typically used. Figure 3-5 shows the number of times that each configuration was used for each 15-minute time period throughout the day. As can be seen, configuration 22|13 was used twice as much as any other between 1600 and 2200 hours. This is likely because of the typical wind direction at this time of day.

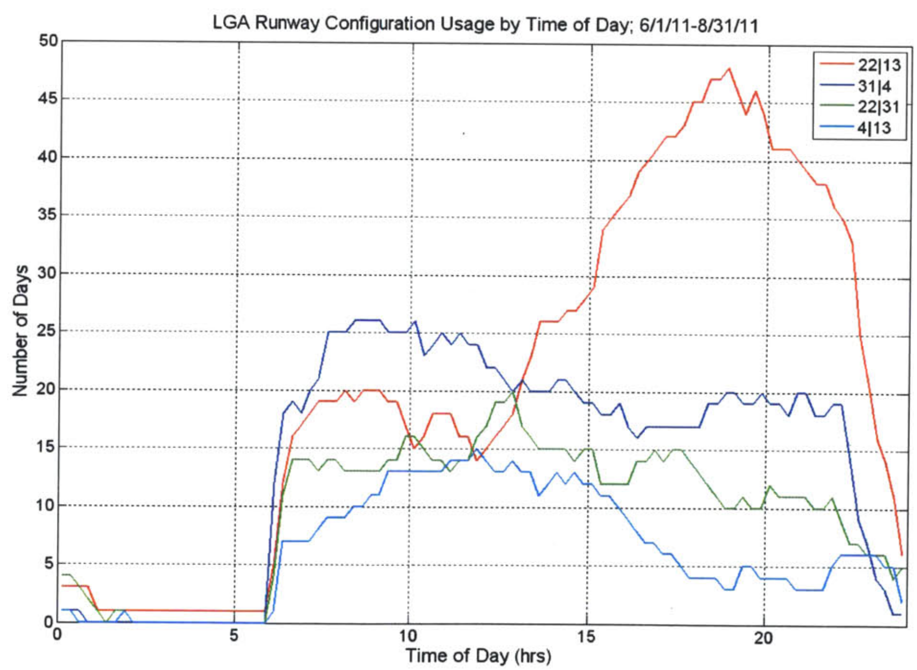

Figure 3-5: LGA runway configuration usage by time of day; Summer 2011

3.5 Queue Boxes

To determine how many aircraft are simultaneously in the runway queue, it is necessary to define queue boxes for the runway thresholds. At LGA there are only two runways, each of which has two departure directions, so four different departure runway queues were defined. However, Runway 22 is seldom used for departures, and is not included in the top four configurations being presented in this paper. The queue box for Runway 22 was still defined and metrics were calculated when this runway was in use for the purpose of calculating overall values of taxi times and queue lengths for the airport.

Queue boxes were determined by manually observing airport visualizations over the given time period. The main focus was on the times of day with the highest levels of congestion because these times would have the longest queue lengths. Off-peak times were also observed to ensure the taxiing and queuing behavior remains similar to the peak times. Runways 4 and 31 have fairly simple queuing behavior. Neither runway requires taxiing across an active runway to get from the main terminal to the runway threshold so there are no runway crossing queues to be aware of. Similarly, Taxiway B runs parallel to both Runway 4 and Runway 31, leading up to these two runway thresholds, so the main boxes can be defined around those long straight taxiways.

An aircraft is considered to be active in the queue box if it is within the latitudinal and longitudinal boundaries of the box, if it is a departing aircraft, and if it is departing on the runway for which the queue box is defined. There are two reasons for this definition: First, there are certain portions of taxiways that are included in multiple queue boxes, and an aircraft within such a segment of taxiway should not be double-counted for time in queue. These aircraft should only be counted as being in the queue for the runway on which they will eventually depart. The other reason is that some aircraft that need to hold for some period of time will be sent to a non-congested portion of the airport to wait. These aircraft should not be counted as queuing if they are in a queue box for some other runway. An aircraft is considered to be queuing if it is in the queue box, as defined above, and it is traveling at a velocity less than 10 knots. This is because a queued aircraft will generally be stopped, with a velocity of zero, but it could also be slowly advancing forward one spot each time one aircraft departs. If the aircraft is traveling at a velocity greater than 10 knots, it is defined to not be in a queue because it is moving at too great a velocity. This definition of queuing behavior is used to determine how many aircraft are waiting in the queue at any given time.

For some runways, such as Runway 4, an overflow queue box is also defined. The second queue box only becomes active when the first queue box is full. The overflow queue box is needed for peak times when the level of congestion is so high that the queue extends on the taxiway, well beyond the boundaries of the main queue box. The reason that the second box is only activated when the first is full is to prevent an aircraft being wrongly included in the count of queued aircraft. During an off-peak time, when the queue does not extend outside of the first queue box, aircraft on hold may depart from that runway but be held in the second queue box. In this scenario, the aircraft in the second queue box should not be counted as being in the queue. This would give an artificially high number to the average taxi-out time for the corresponding time period.

3.5.1 Runway 4

Runway 4 has two queues that directly feed into the runway threshold: a primary queue and a secondary queue. The primary queue lines up on the taxiway to the east of Runway 4. This is the main queue where aircraft are served in a first-come first-served basis. There is a secondary queue to the west of Runway 4. This queue rarely has more than a few aircraft in it. These seem to be flights with specific departure times. They are given priority over the main departure queue at the proper times. Getting to the secondary queue requires crossing the active departure runway. The secondary queue is sometimes also used for departures coming from the general aviation terminal, which is also located to the west of Runway 4. This prevents these aircraft from having to cross the active runway to line up in the primary queue. As mentioned above, Runway 4 also has an overflow queue. If the main queue is filled up, the overflow queue is activated and stopped aircraft in the overflow are counted towards the total number of queued aircraft. In this case, the overflow queue box is activated when there are eight aircraft in the main queue box. While the main queue box can hold more aircraft, visualizations show that eight aircraft are waiting before the queue begins to overflow. Sometimes a gap is seen in the main queue in order to provide a space for aircraft to cross the runway. Figure 3-6 shows the main, secondary and overflow queue boxes for Runway 4.

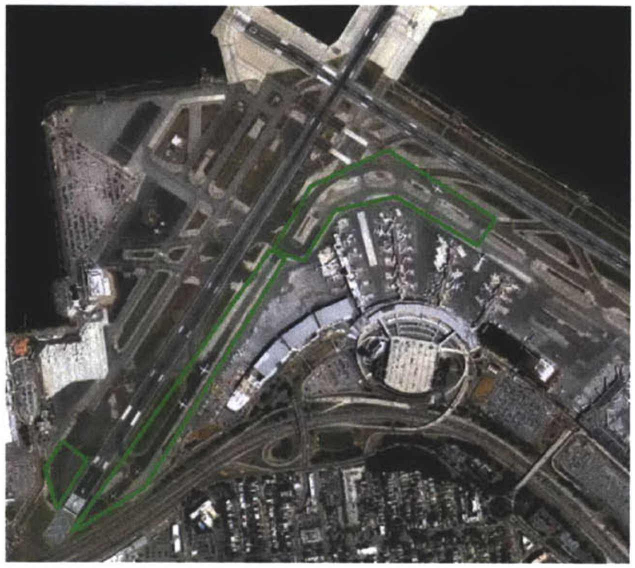

Figure 3-6: LGA queue boxes for Runway 4

3.5.2 Runway 22

As stated above, Runway 22 is rarely used for departures. There is only one taxiway that leads all the way up to the runway threshold, and it runs parallel to the runway, as seen in Figure 3-7. Any queues beyond this length would need to form on the opposite side of Runway 13/31, and follow the taxiway that runs parallel to this runway.

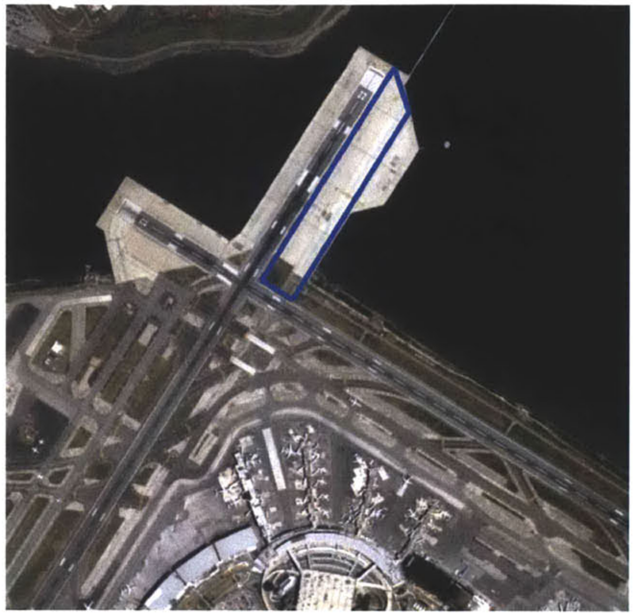

Figure 3-7: LGA queue box for Runway 22

3.5.3 Runway 13

The queuing behavior for Runway 13 is more complicated than the other three runways, mainly because it involves crossing Runway 4/22, but also because there are many taxiways leading to the runway threshold. While taxiing to the threshold of Runway 13 always involves crossing Runway 4/22, there are two different ways of queuing for Runway 13, depending on the arrival runway being used.

In the first case, configuration 22|13, the vast majority of the aircraft are queued on the far side of Runway 4/22, after crossing that runway. In this case, there will only be one or two aircraft waiting to cross the runway, and they will cross as soon as is possible given the traffic requirements of the active runway. For this case, there is one main queue box, which includes four different taxiways. There are three taxiways where aircraft are lined up, all feeding into the runway threshold. Aircraft are lined up based on priority and departure times. When this queue box is full, the overflow queue box on the west side of Runway 4/22 is activated. In this configuration, the crossing of Runway 4/22 takes place closer to the Runway 4 threshold. This configuration is the predominant one for the months studied, and will be discussed in further detail in the next chapter.

In the second case, configuration 4|13, some aircraft are queued up at the runway threshold but the majority of the aircraft are in a queue waiting to cross Runway 4/22. In this case, the controllers seem to intentionally queue aircraft on the opposite side of the runway. It is not merely a matter of waiting for an opportunity to cross the runway. This distinction is important for finding accurate taxi times and time in queue. Typically, the amount of time spent waiting to cross a runway is subtracted from the total time in queue because this amount of time is not spent in a runway threshold queue. If the arrival runway were not in use, this time of waiting would not be necessary. However, in this case, there may be as many as 15 aircraft waiting on the near side of Runway 4/22 waiting to cross. This is a different scenario than in configuration 22|13. Here the runway crossings, rather than occurring down near the Runway 4 threshold, occur right next to the intersection with Runway 13/31. Even though the main queue box is the same as for configuration 22|13, any remaining aircraft are queued up waiting to cross the arrival runway. The time spent in the runway crossing queue should be added to the time in the main queue for the runway threshold to generate the total time in queue.

The main, overflow and runway crossing queue boxes for Runway 13 can be seen in Figure 3-8.

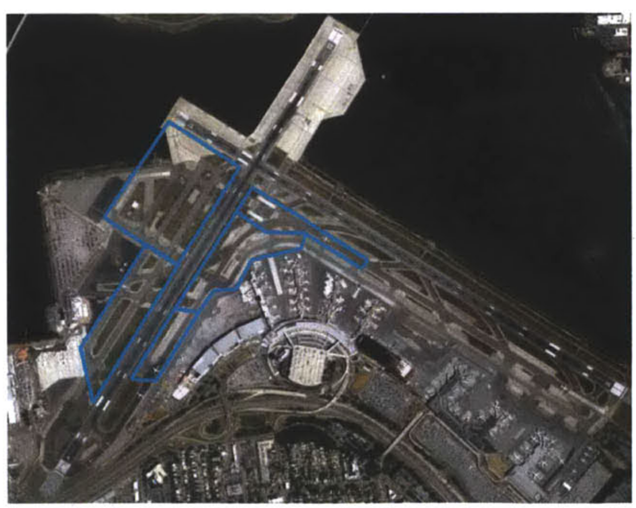

Figure 3-8: LGA queue boxes for Runway 13

3.5.4 Runway 31

Similar to Runway 4, Runway 31 has a long straight taxiway, running parallel to the runway. It does not, however, have a secondary queue on the opposite side of the runway. It has been observed by manually studying the airport visualizations that even the peak-period runway queues do not extend beyond the main queue box so no overflow queue box is necessary. It has also been observed that held aircraft will wait on the secondary taxiway and in the ramp area to the south of Runway 31 shown in Figure 3-9.

Figure 3-9: LGA queue boxes for Runway 31

3.6 Propeller-driven Aircraft

One factor in better understanding factors such as throughput rates and runway capacity is the number of propeller-driven aircraft utilizing the runway. Props can often take off and quickly fan out in a different direction from the previous departure so that the minimum spacing requirements are not violated. Therefore, the time between successive takeoffs is less when a prop trails another aircraft. A significant number of propeller operations at a given airport can lead to changes in the expected throughput of a runway[11].

At LGA, there is not a significant amount of propeller-driven aircraft operations. Of all the departing or arriving aircraft in the three months studied, the average number of props per day was 2.4. The maximum number seen was eight props in a single day over the three summer months of 2011, which does not seem to significantly affect runway operations.

3.7 Overall Analysis

3.7.1 Taxi Times

The analysis of taxi times at an airport is useful to all stakeholders in the system, including the airlines, airports and the FAA. Airlines can use predicted taxi times to better schedule block times for their flights. Airports that are known for large levels of congestion and long taxi times are especially useful to study, because congestion adds large taxi delays on top of airborne time to the total block times. The FAA can use this information for monitoring performance of the entire airspace system, as well as for proposing new policy regulations. Airports can use the information to understand normal operations, and determine how to improve operations in off-nominal cases.

Average taxi-out times in this paper are determined for each 15-minute time interval of the day. Any flights which begin their taxi out during that 15-minute time period are included into the average.The overall analysis includes all configurations, not just the four being studied here. The analysis will be broken down by configuration in the next chapter.

Because not every configuration is being used for a whole day or even for the same time periods over various days, the taxi-out times are summed for all flights that are both in that time period and under that specific runway configuration. This total is then divided by the number of flights falling into that group to calculate the average values.

For the overall analysis, the total taxi-out time is summed over all flights in a given 15-minute time period during the day, over all days studied. Again, the total number of flights falling into this category is counted and used to calculate the average taxi-out times. Two graphical representations of the taxi-out times are shown in Figure 3-10.

The graph on the left in Figure 3-10 shows the average departure taxi times from all configurations, broken down into 15-minute time periods. This includes all departing flights over the 92 days of June through August 2011. There are two peaks which correspond to

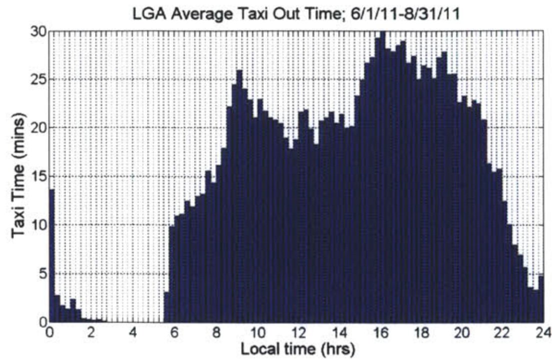

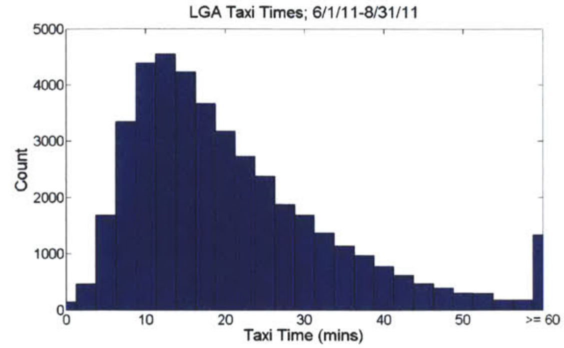

Figure 3-10: (Left) LGA average taxi-out times across all configurations; (Right) LGA distribution of taxi-out times

the morning and evening departure pushes. In the morning, the average taxi-out time peaks at 0900 hours, with an average time of 23 minutes. There is a steady period of longer taxi times in the afternoon to early evening. From 1500 to 2000 hours the average taxi-out times stay above 20 minutes. The peak time for the whole day is at 1600 hours, with an average taxi time of 28 minutes. The times drop steadily after 2100 hours.

The righthand graph of Figure 3-10 shows the breakdown of all the taxi times into bins of 2.5 minutes each. This presents the overall counts of various taxi times. It can be seen that the most frequently observed departure taxi times are between 11.3 and 13.8 minutes long, with over 4500 occurrences over the three months being studied. The calculated average departure taxi time of all flights in the study is 20.3 minutes.

While arrival taxi times tend not to vary as much as departure taxi times, it is still useful to study them to see if the congestion on the ground is affecting the length of the arrival taxi times. Figure 3-11 shows the arrival taxi times for all configurations at LGA. There is little variation in the times throughout the day. The average for all flights is 6.8 minutes.

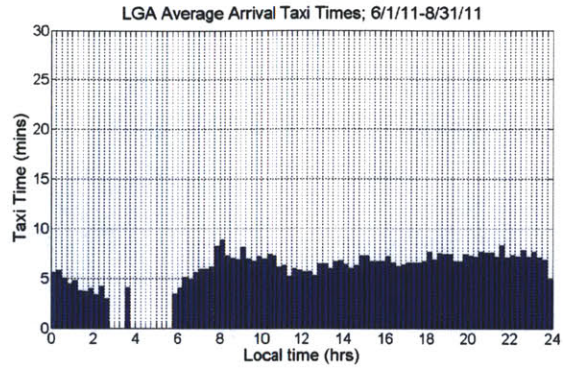

Figure 3-11: LGA average arrival taxi times across all configurations

3.7.2 Arrival and Departure Counts

It is useful to study arrival and departure counts to determine what average throughput is, to come up with a maximum operational capacity and to compare that with the theoretical capacity of the airport. These counts can be compared across multiple configurations to determine which runway configuration has the greatest throughput, and what the variation is between runway configurations. Again, these can be compared against the theoretical throughputs and any significant difference can be analyzed. While it is possible that the capacity estimate is incorrect, operational improvements may also be able to increase the observed throughput. It is therefore helpful to not only calculate the departure and arrival counts, but also to study the airport visualizations in the given runway configuration. These clearly show the interaction of departures, arrivals and the departure runway queues.

For this study, arrivals and departures have been measured in terms of counts per 15-minute time period. The sum of all departures during a given time period is divided by the number of days to determine the average number of departures per time period. The same method is used for the average number of arrivals.

As would be expected from a slot-controlled airport, the departure and arrival counts stay fairly constant throughout the day as compared to some other unregulated airports. As can be see in Figure 3-12, departures begin at 0600 hours and quickly taper off after 2200 hours. The maximum average number for a time period is just over eight aircraft per 15 minutes and occurs at 1215 hours. The average number of departures slowly declines over the course of the day until reaching a rate of five aircraft per 15 minutes at 2200 hours.

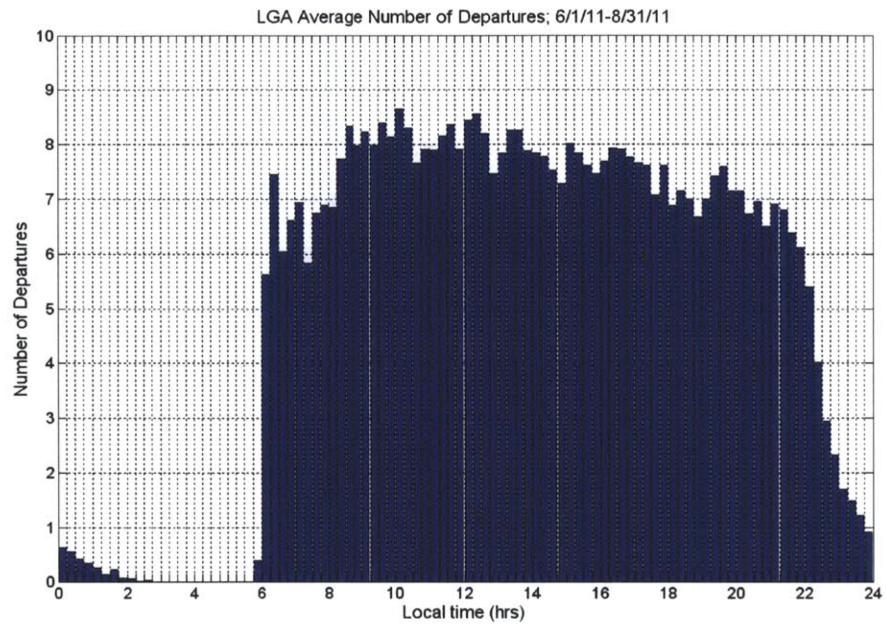

Figure 3-12: LGA average number of departures across all configurations

The average arrival behavior over all the runway configurations is shown in Figure 3-13. The arrivals ramp up shortly after 0600 hours and reach a maximum at 1030 hours. The counts drop off after 2330 hours, although arrivals can be seen until 0200 hours. The peak average arrival count is seen at 1045 hours, with just over nine arrivals per 15 minutes.

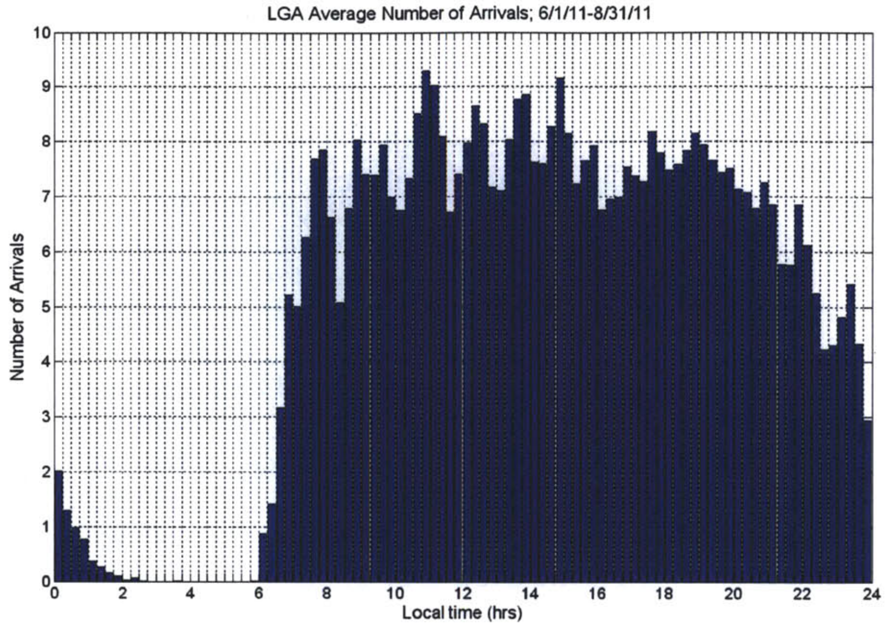

Figure 3-13: LGA average number of arrivals across all configurations

3.7.3 Departure Rates

The departure rate is defined to be the number of takeoffs on a given runway under a certain runway configuration. In this study, the departure rate will be measured as the number of departures per 1-minute time interval. The departure rate for the same departure runway may be different under different configurations. For example, it is expected that the departure throughput for configuration 31|31 would be less than for 22|31. Even though both configurations are using Runway 31 for departures, the former will have fewer departures because the same runway is being used for both arrivals and departures, and can therefore serve fewer departures in the same amount of time.

It also matters whether two intersecting runways are being used for departures and arrivals at the near end of the intersection, or the far end. If the intersection is at the near end, a departing or arriving aircraft will clear the intersection quickly, and the next operation on the opposite runway can take place immediately. If at the far end, there is a longer takeoff roll or landing roll time that must be factored in before the first operation has cleared that intersection, and the successive operation can take place.

Departure rates will be discussed in detail for each configuration at LGA in the following chapter. The rate is represented as a function of the number (N) of departing aircraft on the surface. The number of departing aircraft includes all active departures, both taxiing aircraft and those waiting in queue for the runway. Departure rate curves generally increase up to the capacity of the runway, hold steady there to some level of N and then begin to decrease. This will be discussed further in the next section.

A configuration-dependent departure rate curve is more useful than an overall one for the airport across all configurations, because the rates themselves depend significantly on the configuration in use. Each one of these curves will be presented individually in the next chapter.

3.7.4 Saturation and Throughput

For studies of congestion, taxi times and throughput, one useful measure is finding the point at which the airport saturates. When there is not much traffic on the ground and no queue for the runway, the number of departures is solely based on the demand for the runway. As the demand increases, so does the number of departures. At some point, the airport reaches its maximum capacity, beyond which an increase in demand will not increase the number of departures because the runway can only accommodate a limited number of departures per 15-minute time period. However, beyond this point, if the number of aircraft on the ground increases, there may be a decrease in runway utilization. As the airport taxiways become saturated, the congestion limits the movement of aircraft on the surface and may actually hinder runway utilization. Generally arrival taxi times are considered to stay fairly constant because they are not waiting in a queue for the runway. However, as the congestion on the ground increases, the arrival aircraft have a difficult time reaching their gates. They may have to take a much longer taxi route or may even have to wait for some period of time in a queue with departing aircraft just to be able to reach their intended gate.

For the above reasons, it is useful to study saturation points at an airport. The saturation point is defined to be the number of total aircraft on the ground beyond which there is no increase in departure throughput, and there may even be a decrease in throughput.

Throughput curves are presented for each of the four main configurations in the following chapter. From these curves, a saturation point can be determined. The curves will show the departure rates parameterized according to the number of arrivals.

3.8 Conclusions

The overall study of LaGuardia Airport through analysis of ASDE-X data has shown some consistent patterns in surface operations. In the next chapter, each configuration will be explored individually to see how much these patterns differ from one configuration to another.

The average numbers of departures and arrivals tend to be fairly constant throughout the day, with departures maxing out at an average of 8.5 per 15 minutes, and arrivals reaching a maximum average rate of 9.4 per 15 minutes. Despite a relatively constant departure rate, the demand for departures varies within the day, with a major push in both the morning and evening. This leads to long queues and increased taxi times in the morning around 0900 hours and for a period of time in the late afternoon and evening. On the other hand, these changes in departure demand do not appear to interfere with arrivals, as arrival taxi times stay constant throughout the day.

Chapter 4

Analysis of LGA Configurations

4.1 Introduction

As stated in the previous chapter, there were four main runway configurations in use at LaGuardia Airport over the time period being studied, June through August of 2011. In decreasing order of use they are: configuration 22|13 (used 37% of the time), configuration 31|4 (26% usage), configuration 22|31 (17%), and configuration 4|13 (used 12% of the time). Together, these configurations are used 92% of the time periods. By using ASDE-X data to study these four configurations, the vast majority of operations at the airport can be categorized and analyzed.

4.2 Motivation

The previous chapter used ASDE-X data as a tool to broadly describe LGA airport operations. Average values were calculated to show what typical operations at the airport are like, without regards to any other variables. Some other variables that can significantly influence airport operations are runway configuration, weather, and day of the week.

Weather probably has the biggest impact on airport operations, since it can shut down an entire airport leaving hundreds of passengers stranded and trying to rebook flights. Even without any significant weather at the given airport, adverse weather in the surrounding airspace can affect surface operations. Flights may be delayed for a specific amount of time or indefinitely, since weather conditions are unpredictable. For aircraft taxiing, this may mean holding on the ground for long periods of time. Controllers may send these aircraft to a specific location for holding, and some aircraft may even return to the gate because of the 3-hour tarmac delay rule. While weather will not be explicitly taken into account in this thesis, a future application of ASDE-X data would be merging it with weather forecasts to further classify airport operations and make predictions for the future using the merged data.

Airport operations also tend to vary based on the day of the week, depending on the market (business vs. leisure) that the airport mostly serves. One classification to look at is the difference between weekday and weekend operations. The metrics may not vary much throughout the weekdays, but often for passenger operations, the weekends will have lower demands and shorter taxi-out times.

Finally, many of the metrics discussed in the previous chapter will vary depending on the runway configuration. The configuration has a big impact on runway throughput. Taxi-out times will vary both because of differences in throughput, but also based on how close the departure runway is to the starting location of departing aircraft. The four main runway configurations at LGA will be discussed below. The average values for each configuration will be presented and then reasons for the differences between them will be discussed.

4.3 Configuration 22|13

The most commonly used runway configuration has arrivals on Runway 22 and departures on Runway 13. The sections below will discuss queuing behavior, taxiing, arrivals, departures and throughput.

4.3.1 Queuing Behavior

Based on observing the visualizations of the ASDE-X data, Runway 13 exhibits different queuing behavior depending on the arrival runway. When Runway 22 is the arrival runway, the aircraft first line up on the three taxiways in the main queue box. After that has filled, typically aircraft will queue up on the opposite side of Runway 22 in a line on the taxiway running parallel to Runway 13. Figure 4-1 shows an example of this.

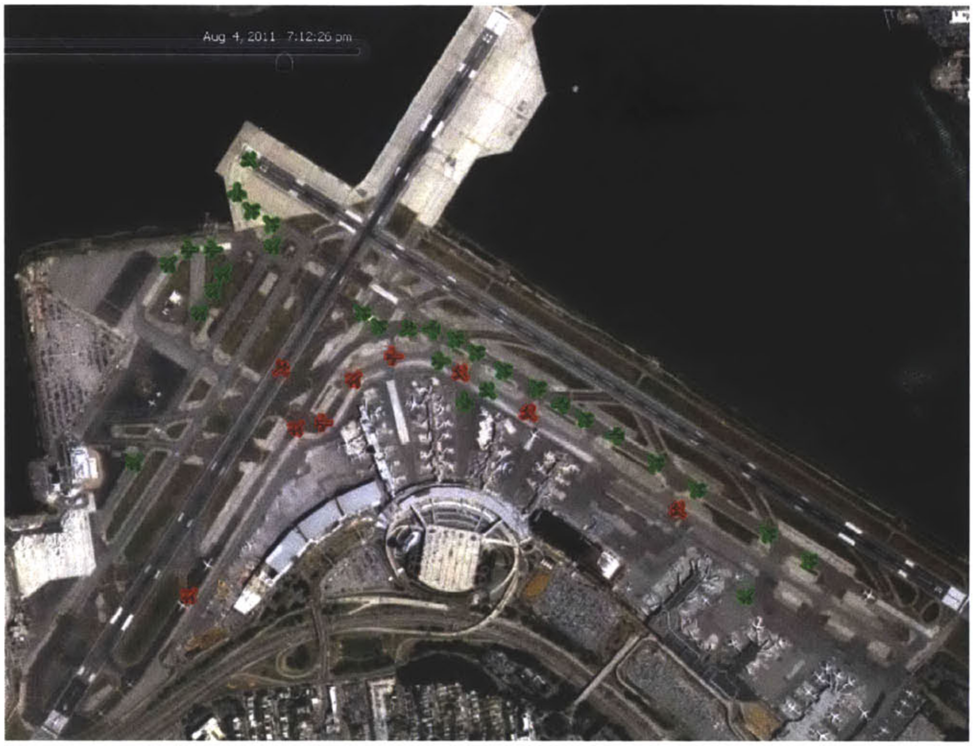

Figure 4-1: Queue for configuration 22|13

4.3.2 Taxi Times

As with the overall departure taxi times, configuration 22|13 exhibits two peaks in taxi-out times corresponding to the morning and evening pushes. This can be seen in Figure 4-2. The morning peak is between 0900 and 1100 hours with a maximum value of 25 minutes. There is a second peak between 1700 and 2000 hours. The maximum average taxi time for a 15-minute time period is 30 minutes. The overall average taxi time of all aircraft is 24.3 minutes.

The arrival taxi times for configuration 22|13 do not vary much from the overall arrival taxi times. The average time is 7.1 minutes.

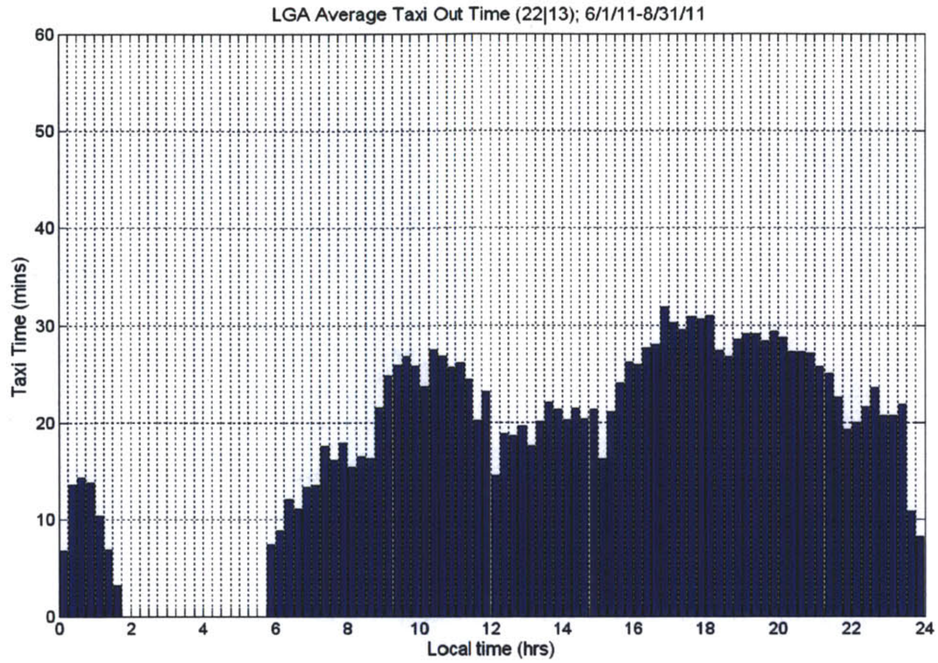

Figure 4-2: Average taxi-out times for configuration 22|13

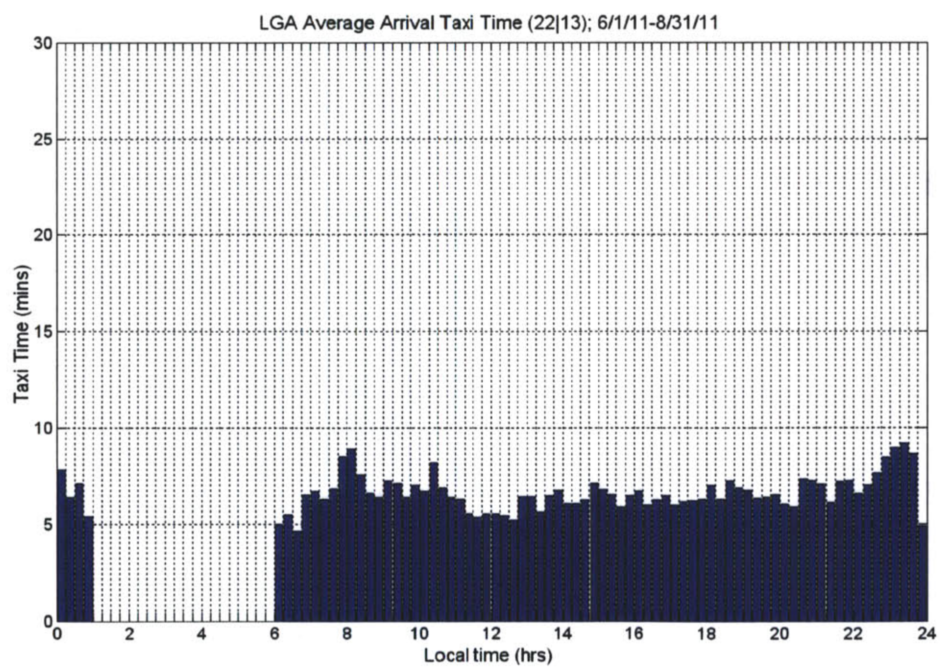

Figure 4-3: Average arrival taxi times for configuration 22|13

4.3.3 Departure Rates

The analysis of departure rates links the number of departures to the number of active departing aircraft on the ground. Figure 4-4 shows the departure rates for configuration 22|13. In this chart, *N* is the number of departures that have already pushed back and are taxiing to, or waiting in queue for, the departure runway. The top graph shows the departure rate as a function of *N*. The bottom graph shows the number of observations for each value of *N*. The value of *N* is measured every minute, and each one-minute interval with a certain value of *N* is summed up.

For this configuration, the maximum departure rate is 0.6, which is equivalent to 36 departures per hour. The departure rate holds fairly constant around 0.6 aircraft per minute for values of N between 10 and 27. When N exceeds 27, the departure rate dips down to 0.46 aircraft per minute. The most common value of N is 8.

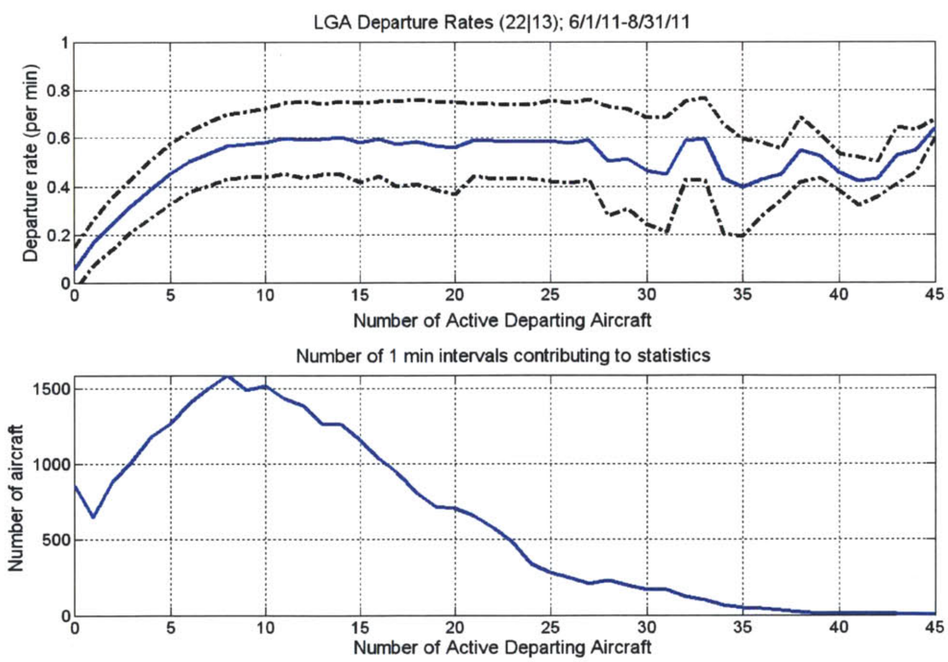

Figure 4-4: Departure rates for configuration 22|13

Figure 4-5 shows how the number of active departing aircraft (*N*) varies by time of day. There are peaks at 0900 and 1700 hours, which correspond to the peaks in taxi times during the day.

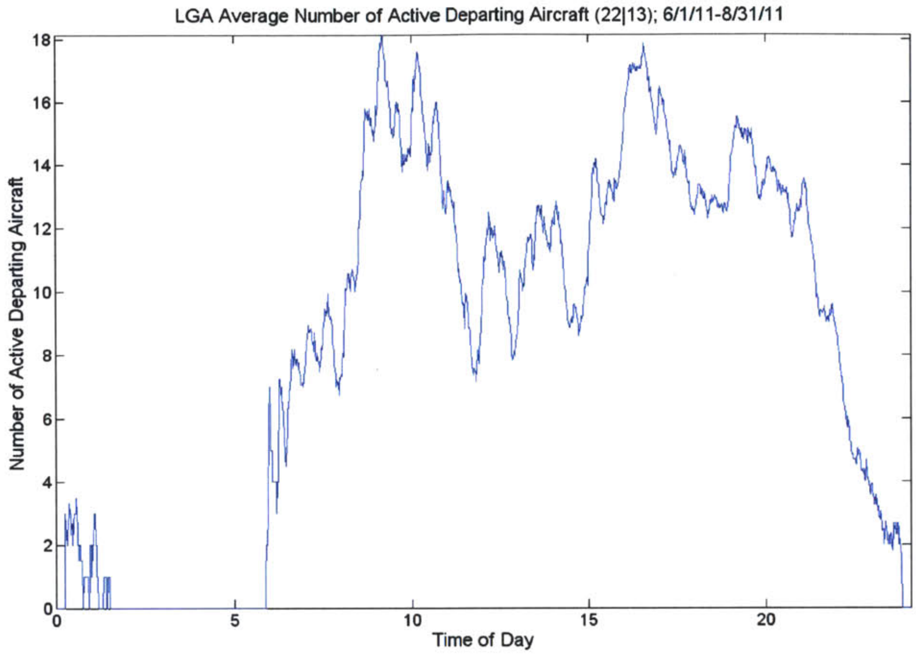

Figure 4-5: Departing aircraft counts for configuration 22|13

4.3.4 Throughput

The throughput parameterized by number of arrivals is shown in Figure 4-6. For configuration 22|13, the departure throughput does not vary with the number of arrivals. This is because the runway intersection is close to both the departure threshold and arrival threshold. An aircraft waiting to depart only needs to wait the amount of time that it takes for an arrival to cross the intersection, before it can begin its takeoff roll. The maximum throughput for all levels of arrivals is 8.7 departures per 15 minutes. Once the number of active departures reaches 11, the number of departures holds constant at the maximum value.

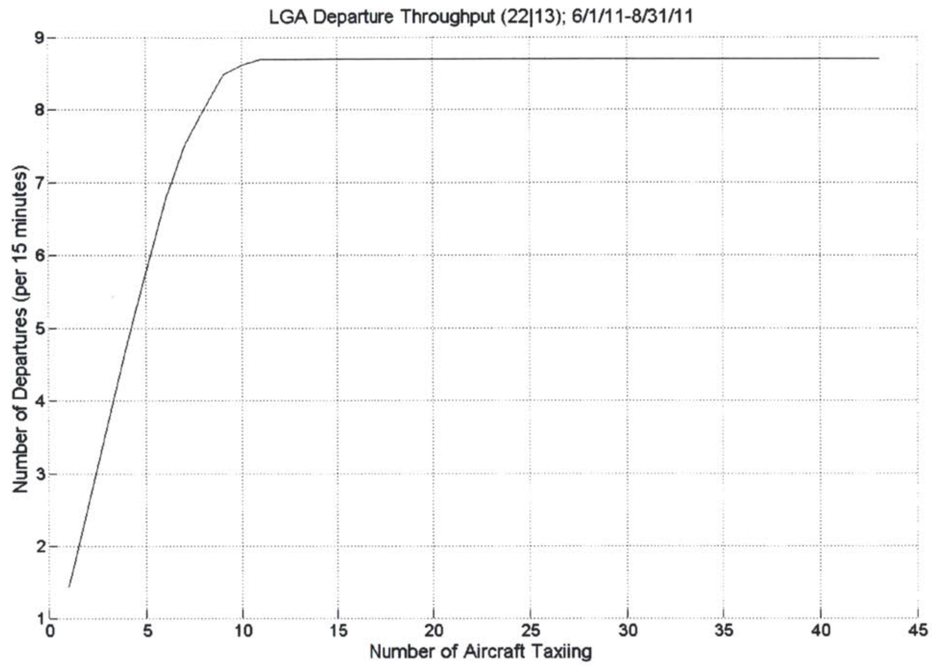

Figure 4-6: Throughput parameterized by arrivals for configuration 22|13

4.4 Configuration 31|4

The second most-used runway configuration has arrivals on Runway 31 and departures on Runway 4. The same metrics mentioned above will be discussed for this configuration.

4.4.1 Queuing Behavior

As stated in the previous section, Runway 4 has a main queue and a secondary queue. The secondary queue appears to be EDCTs, or aircraft coming from the Marine Air Terminal to avoid crossing the departure runway. When the first queue box of the main queue fills up, an overflow queue box becomes active, and any stopped aircraft in that box are also counted as being part of the queue. An example of such a queue is shown in Figure 4-7.

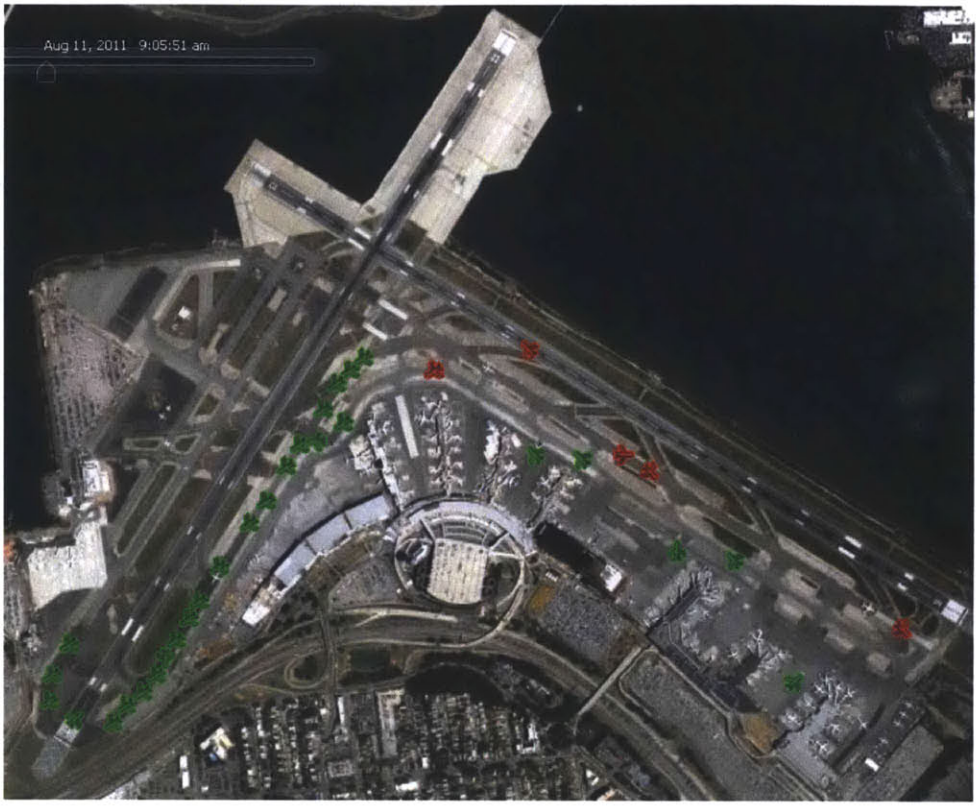

Figure 4-7: Queue for configuration 31|4

4.4.2 Taxi Times

The average departure taxi times for configuration 31|4 are shown in Figure 4-8. This configuration also has a peak in the morning around 0900 hours, with an average taxi time of 27 minutes. However, this configuration does not have the same increase in values over the afternoon and evening periods. The average taxi time across all departures in this configuration is 21.9 minutes.

The arrival taxi times are shown in Figure 4-9. The average taxi time for all arrivals under this configuration was 6.1 minutes.

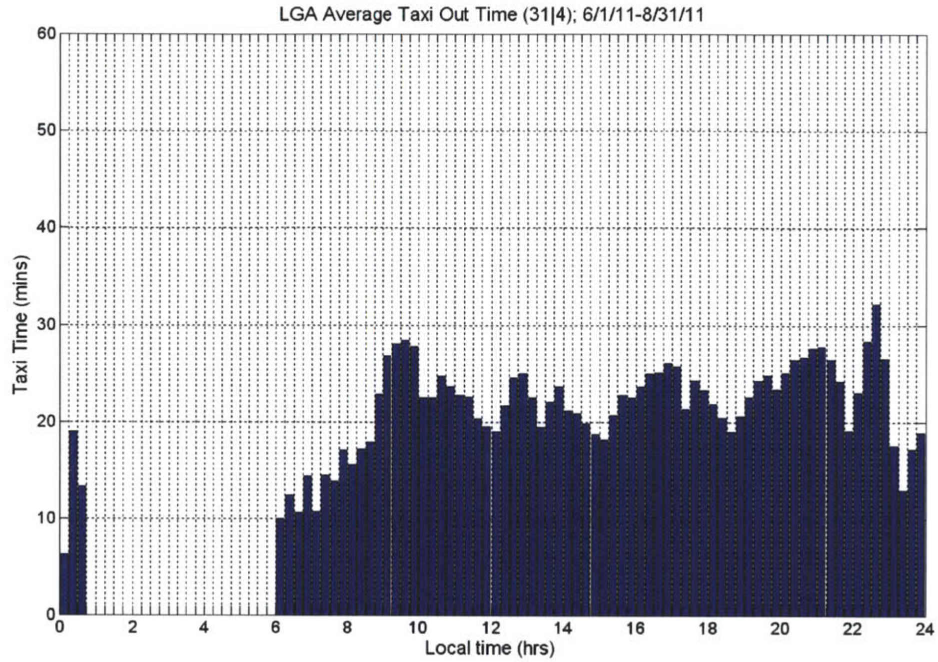

Figure 4-8: Average taxi-out times for configuration 31|4

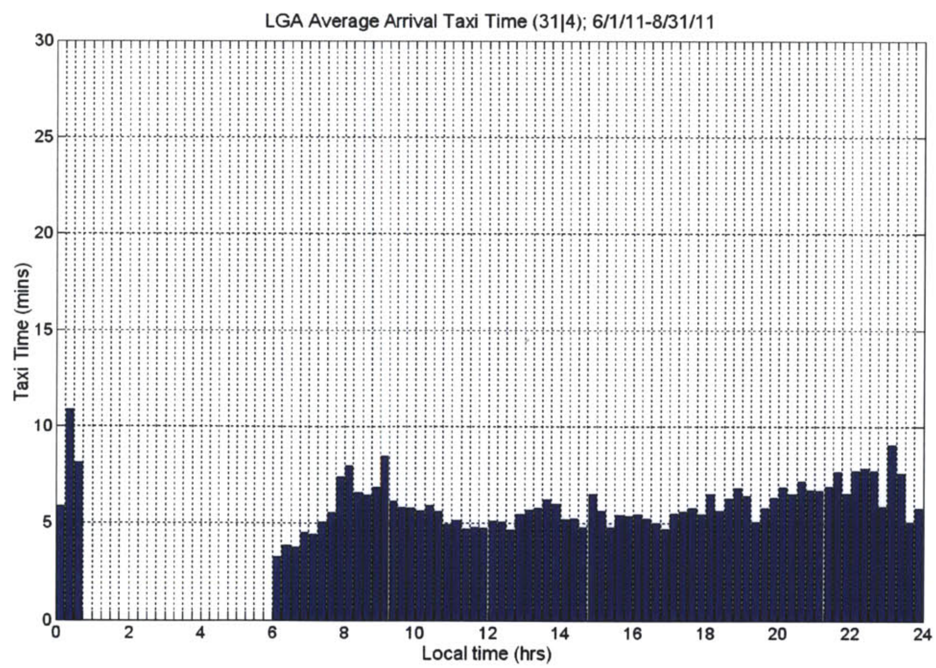

Figure 4-9: Average arrival taxi times for configuration 31|4

4.4.3 Departure Rates

The departure rates for configuration 31|4 are shown in Figure 4-10. The rate increases steadily until N reaches 15, with a rate just over 0.6 departures per minute. Beyond that the rate continues gradually increasing until N reaches 28 and a rate of 0.67 departures per minute. The most common value of N seen is 10.

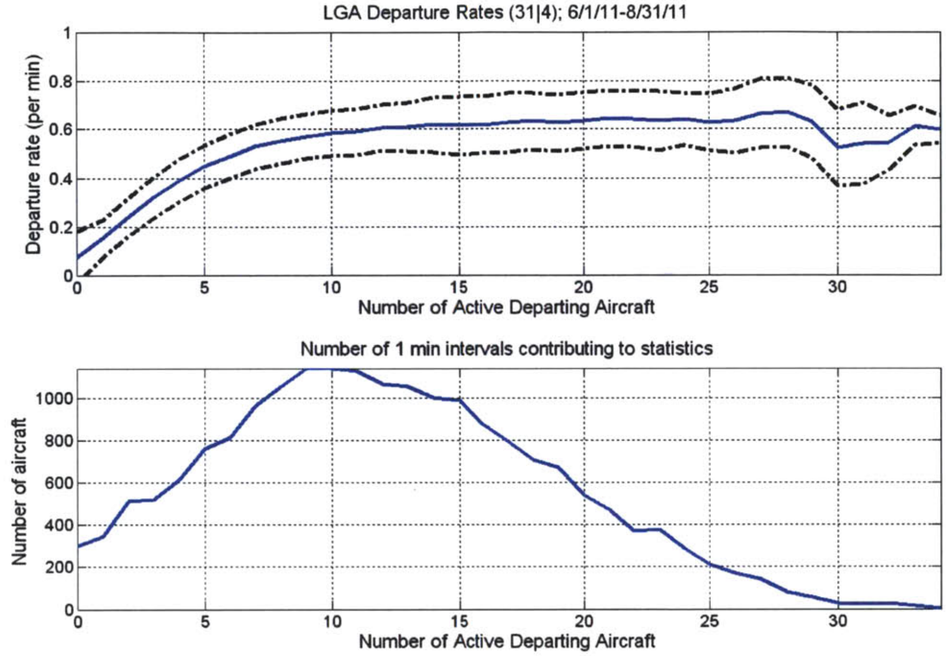

Figure 4-10: Departure rates for configuration 31|4

The variation in the number of active departing aircraft throughout the day is shown in Figure 4-11. The shape of the curve closely parallels the rise and fall of departure taxi times. The first peak in traffic is around 0900 hours with another peak in the afternoon around 1600 hours.

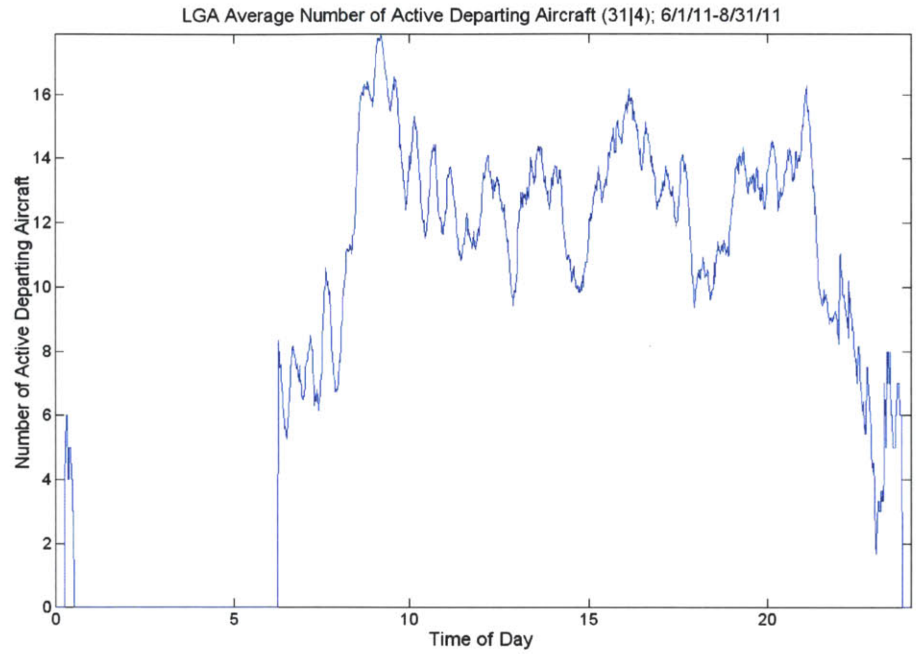

Figure 4-11: Departing aircraft counts for configuration 31|4

4.4.4 Throughput

The throughput for configuration 31|4 is the one that varies the most by the number of arrivals. This is likely because the intersection of the two runways is at the far end of both the arrival and departure runways. When there are no arrivals, the throughput reaches a maximum rate of 12.6 departures per 15 minutes as N increases. The bottom of the curve is the rate when there are 16 arrivals. In this case the throughput reaches a maximum of 7.7 aircraft per 15 minutes. The increase in arrival levels causes a decrease in departure throughput of 5 aircraft per 15 minutes. In contrast, configuration 22|13 showed virtually no difference in departure throughput based on arrivals (Figure 4-6).

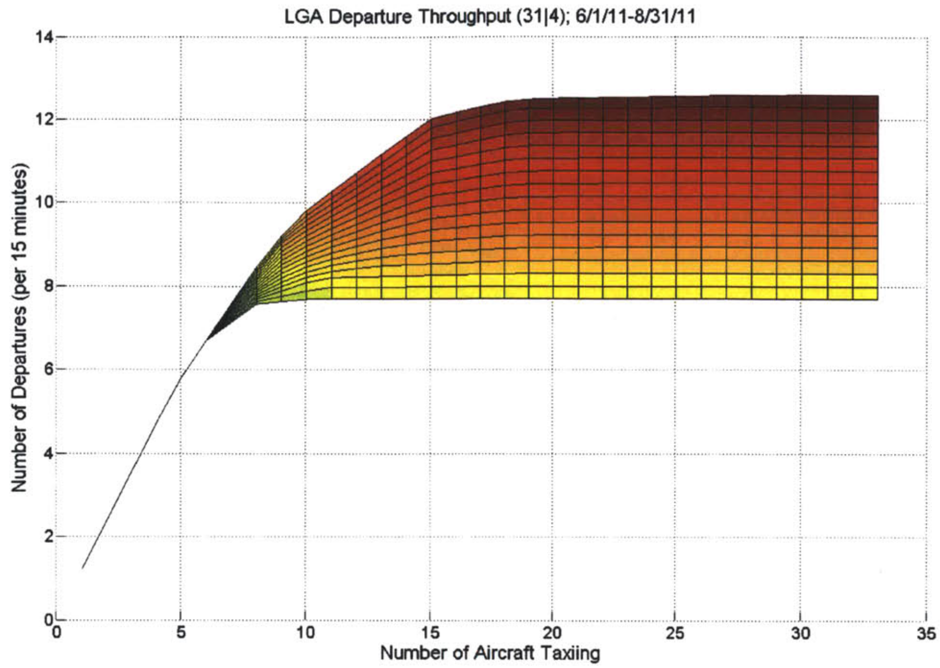

Figure 4-12: Throughput parameterized by arrivals for configuration 31|4

4.5 Configuration 22|31

Configuration 22|31 was used 17% of the time in summer 2011.

4.5.1 Queuing Behavior

The queue for Runway 31 follows the taxiway immediately next to the runway. Figure 4-13 shows an example queue for Runway 31 in this configuration. Aircraft that are holding are observed to be waiting in the ramp area just below the runway threshold. Aircraft with EDCTs can be observed to be waiting on the neighboring taxiway, and are given priority over the main queue at their proper departure times.

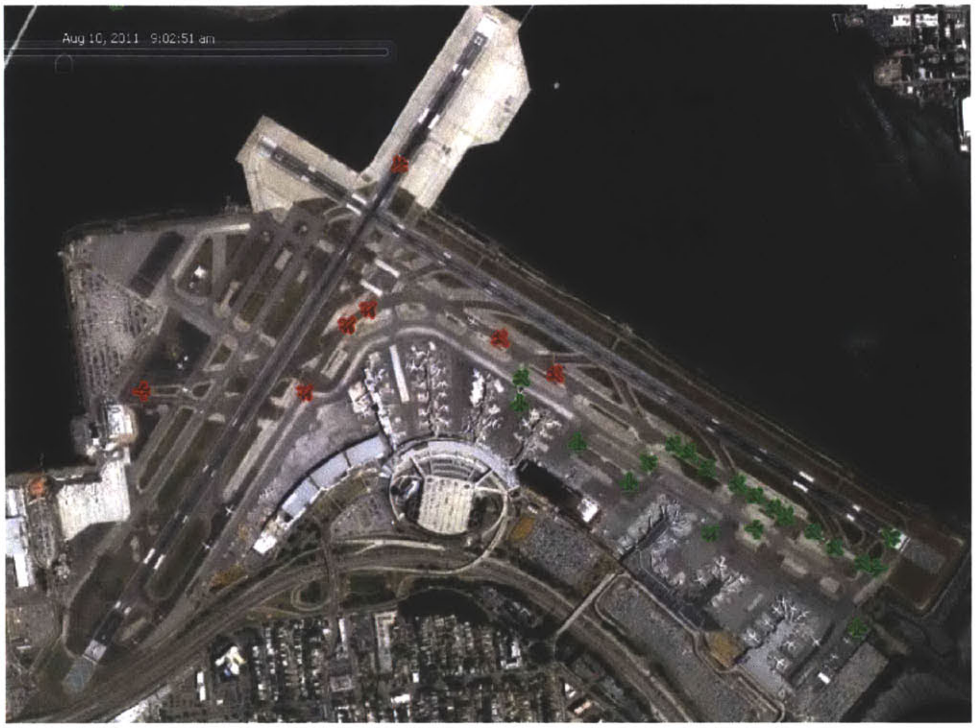

Figure 4-13: Queue for configuration 22|31

4.5.2 Taxi Times

The departure taxi times for configuration 22|31 are lower than for the previous two configurations discussed above. Figure 4-14 shows the average taxi-out times throughout the day in 15-minute time intervals. The maximum average value for any 15-minute time period is an average length of 18 minutes at 0945 hours. The average taxi time for all flights seen in this configuration is 14.5 minutes.

The average arrival taxi time graph, seen in Figure 4-15, looks similar to the other two configurations previously discussed. The average value for all flights is 7.7 minutes.

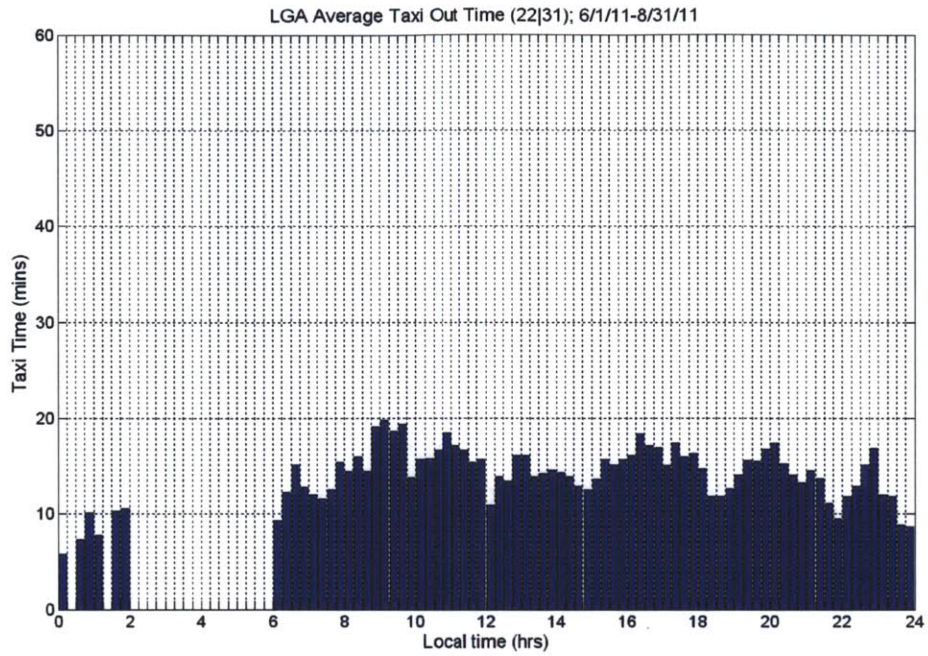

Figure 4-14: Average taxi-out times for configuration 22|31

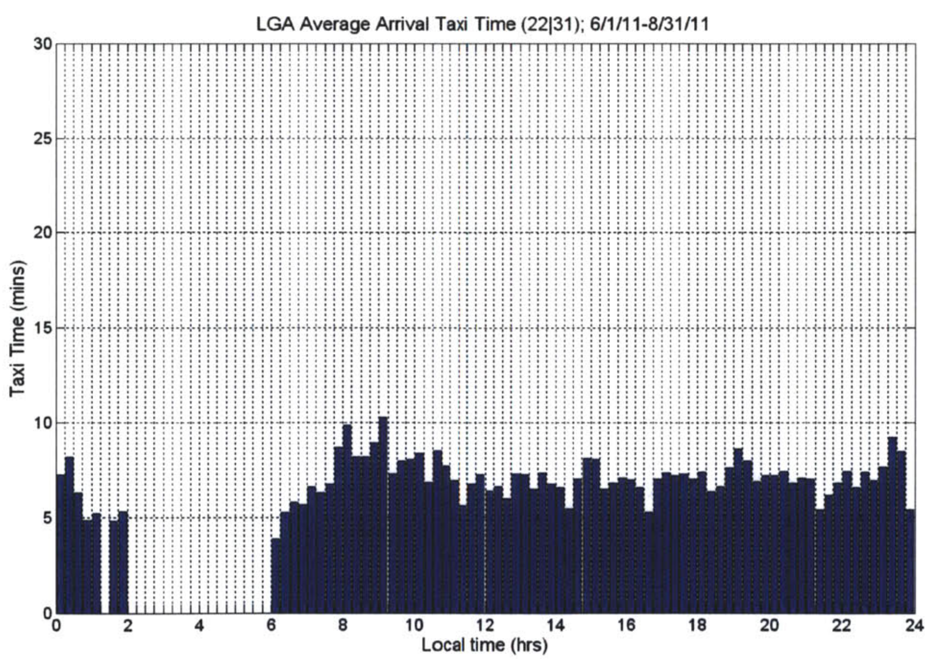

Figure 4-15: Average arrival taxi times for configuration 22|31

4.5.3 Departure Rates

The departure rate for configuration climbs to a value of 0.6 when there are 8 active departing aircraft and then continues gradually climbing to a rate of 0.7 when there are 19 active departures. The most common value of N was 8.

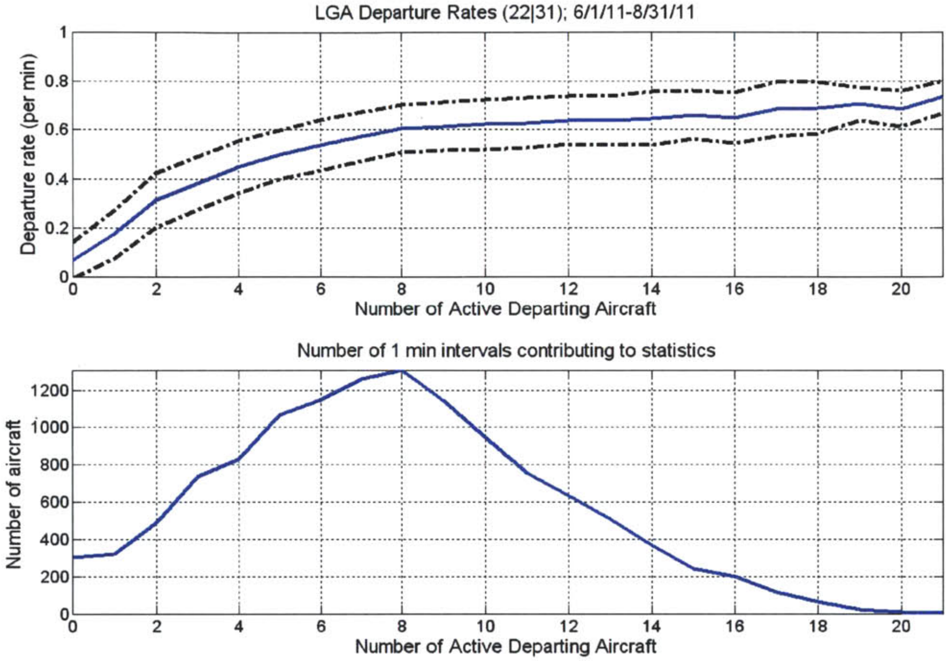

Figure 4-16: Departure rates for configuration 22|31

Again, it is seen that the number of active departing aircraft corresponds to the taxi times seen throughout the day. The maximum averaged number of departing aircraft occurs at 0900 hours with a value of 12. This can be seen in Figure 4-17.

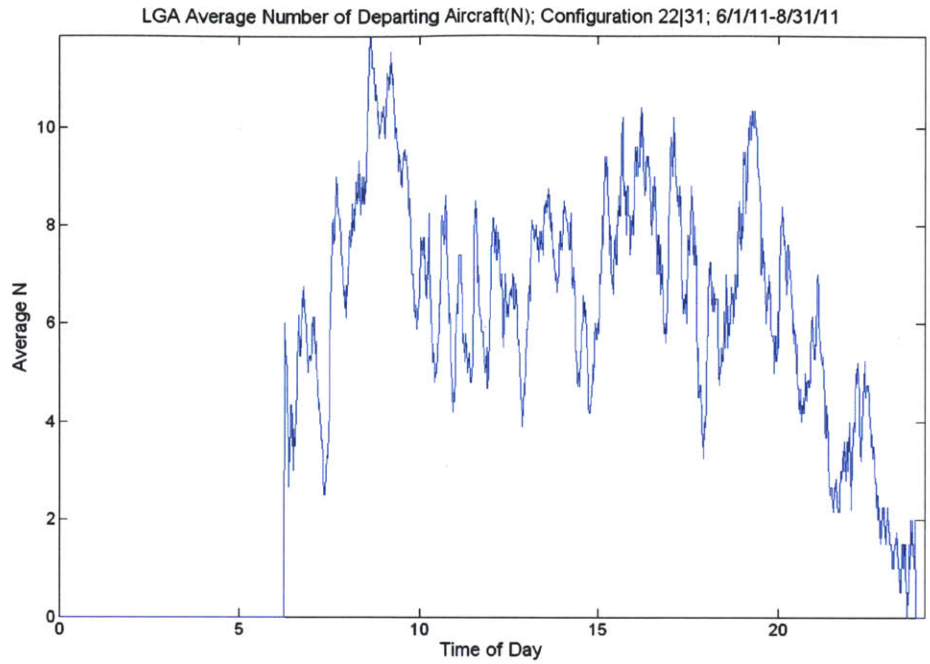

Figure 4-17: Departing aircraft counts for configuration 22|31

4.5.4 Throughput

The average throughput for configuration 22|31 is shown in Figure 4-18, parameterized by the number of arrivals in a 15-minute time period. The maximum rate, when there are no arrivals, is 10.7 departures per 15 minutes. When the arrivals increase to 12, the throughput falls to 9.5 departures. When there are fewer than 10 active departing aircraft, the departure rate does not vary with the number of departures.

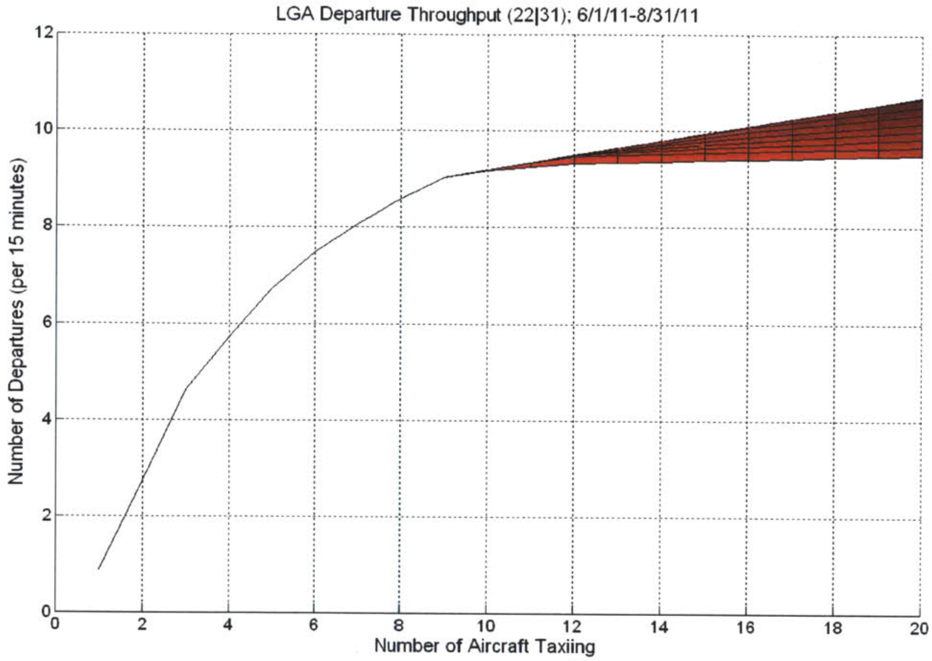

Figure 4-18: Throughput parameterized by arrivals for configuration 22|31

4.6 Configuration 4|13

The final configuration to be presented has departures on Runway 13 and arrivals on Runway 4.

4.6.1 Queuing Behavior

The queuing behavior for this configuration is very similar to configuration 22|13 because they both use Runway 13 for departures. The difference is in the runway crossing location. The location of the runway crossing shown in Figure 4-19 is a common location for this configuration. The visualizations also show runway crossings down closer to the arrival threshold of Runway 4, usually when the queue has filled the main queue box and extended into the overflow queue box.

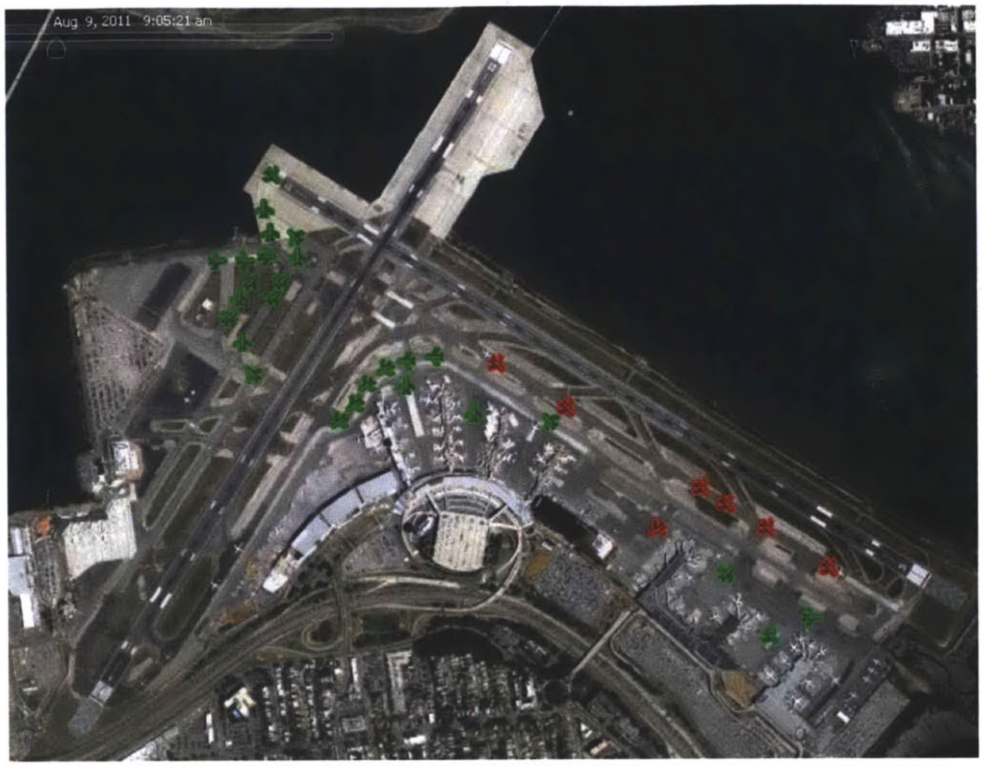

Figure 4-19: Queue for configuration 4|13

4.6.2 Taxi Times

This runway configuration exhibits the highest average departure taxi times. The average taxi-out times exceed 30 minutes at 0900 hours and mostly stay above a half hour the rest of the day, apart from a dip around noon. The maximum value is 57 minutes at 1800 hours. The average taxi time for all flights in this configuration is 31.3 minutes.

The arrival taxi times for this configuration, shown in Figure 4-21, do have some spikes that are not exhibited in other configurations. This is more likely to be due to a small number of data points occurring on days with inclement weather, than any real difference in the configuration. The average for all flights is 7.6 minutes.

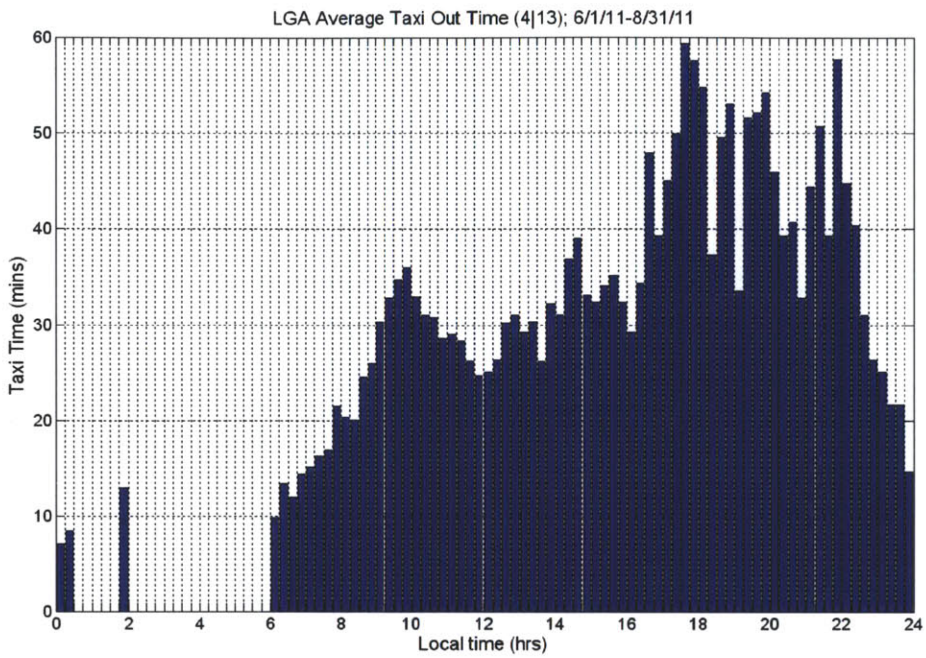

Figure 4-20: Average taxi-out times for configuration 4|13

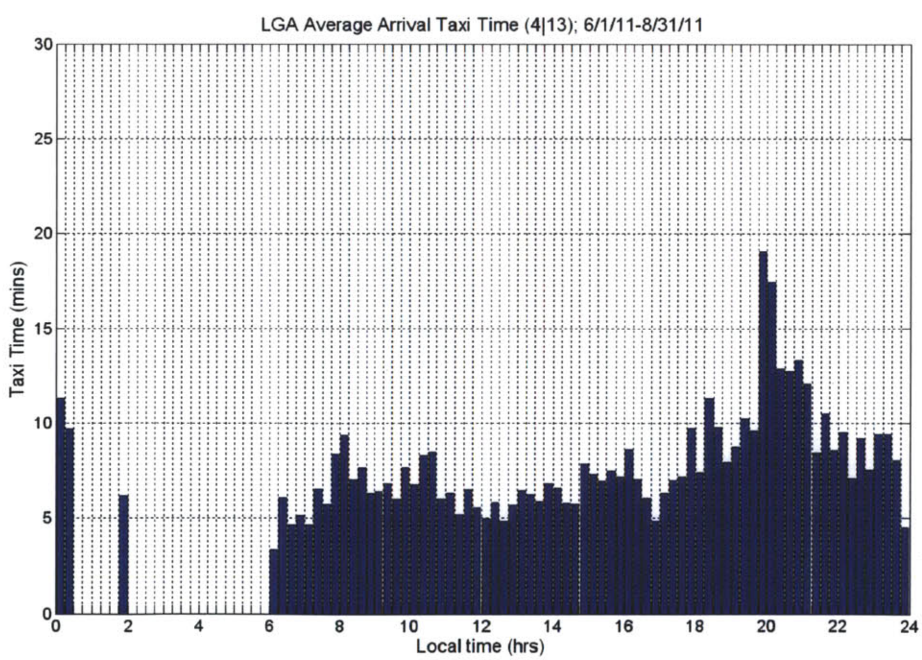

Figure 4-21: Average arrival taxi times for configuration 4|13

4.6.3 Departure Rates

The departure rate for configuration 4|13 steadily increases until the number of active departures equals 15. At this point the departure rate is just under 0.6 aircraft per minute and stays around this level until there are 22 active departures. Then the departure rate dips down to 0.54.

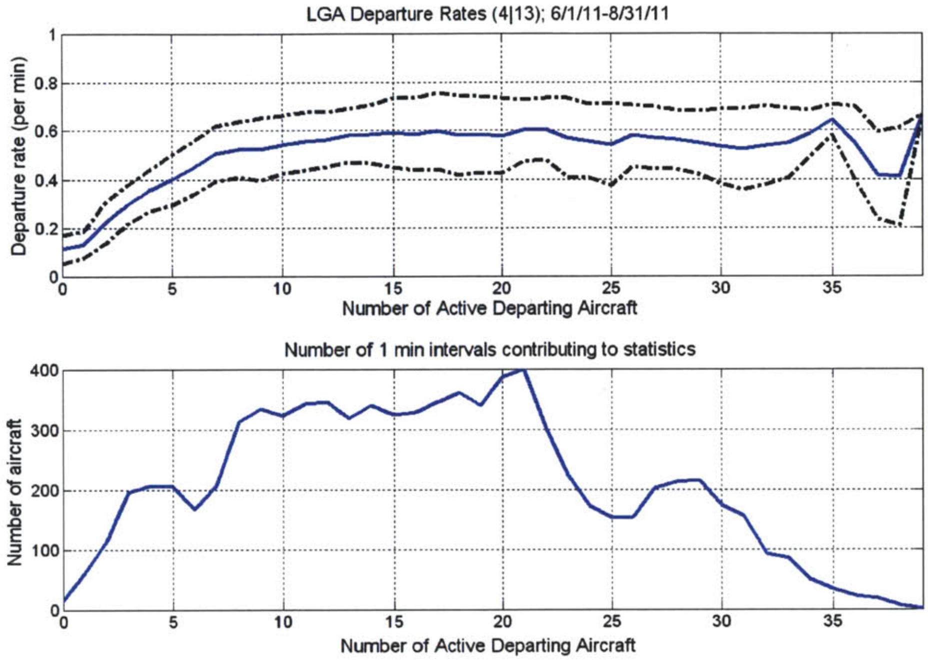

Figure 4-22: Departure rates for configuration 4|13

This configuration has the highest number of active departures throughout the day. This may again be due to the limited amount of data for this configuration.

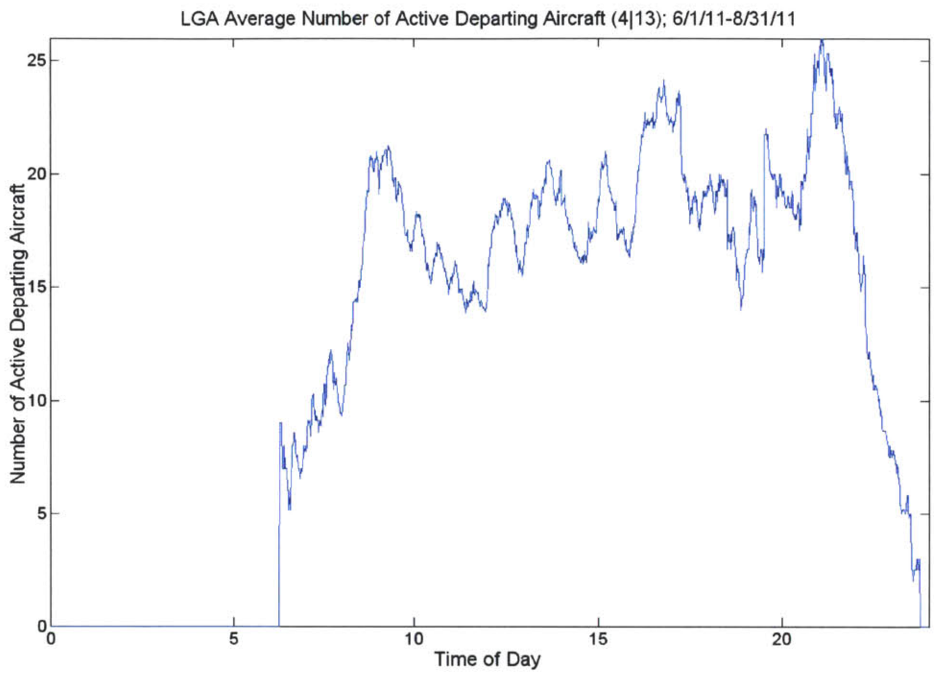

Figure 4-23: Departing aircraft counts for configuration 4|13

4.6.4 Throughput

The departure throughput plot for configuration 4|13, as seen in Figure 4-24, is quite similar to the one for configuration 22|13. The maximum rate is 8.7 departures per 15 minutes when there are no arrivals. There is a slight amount of variation based on the number of arrivals. The rate decreases to 8.5 when there are 12 arrivals.

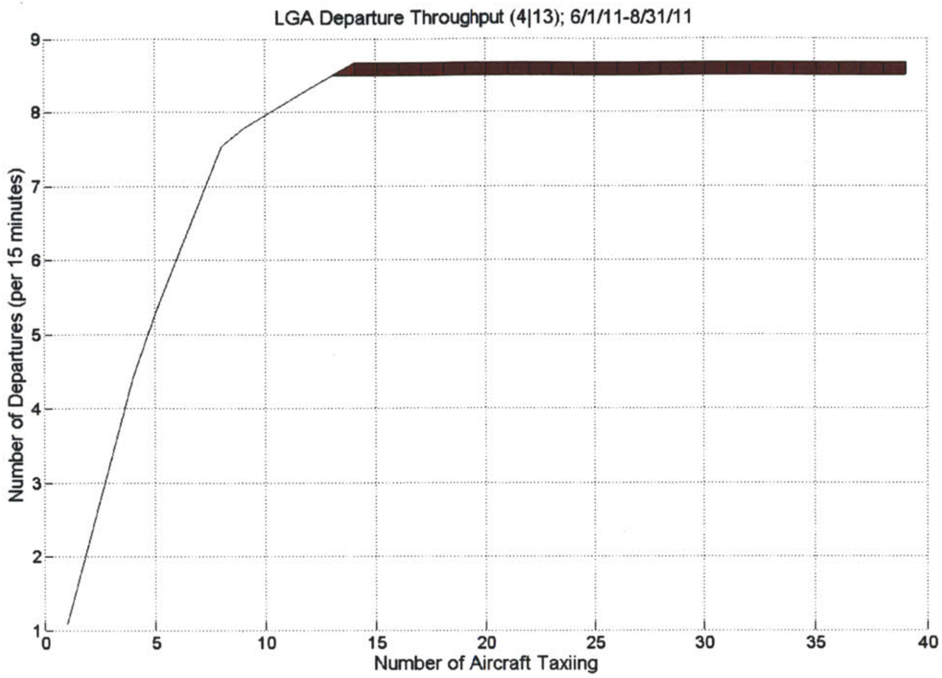

Figure 4-24: Throughput parameterized by arrivals for configuration 4|13

4.7 Analysis

A summary of the statistics for each configuration is shown in Table 4.1.

The two configurations with departures on Runway 13, namely 22|13 and 4|13, are the ones that require departures to cross the arrival runway to reach the departure runway threshold. This makes the taxi behavior a little less straightforward, but also provides the aircraft with more room for queuing on the other side of Runway 4/22. The location of holding aircraft varies with the configuration, depending on where there is room on the airport surface.

Each configuration shows peaks in taxi times around 0900 hours and also between 1700 and 1800 hours. Some peaks are more pronounced than others, but they all show similar characteristics. The configuration with the lowest average taxi time was 22|31 with a length of 14.5 minutes. The maximum average taxi-out time was 31.3 minutes for configuration 4|13, which is more than double the minimum average taxi-out time. This metric varies widely between configurations. The average arrival taxi time over a 15-minute time period were fairly constant, all day long, and across all configurations. The minimum average

Table 4.1: Summary of statistics by configuration

	22|13	31|4	22|31	4|13
Usage	37%	26%	17%	12%
Average Taxi-out Time (min)	24.3	21.9	14.5	31.3
Max 15-min Average Taxi-out Time (min)	30	27	18	57
Average Arrival Taxi Time (min)	7.1	6.1	7.7	7.6
Saturation Point (aircraft)	28	28	19	22
Max Departure Rate (per min)	0.6	0.67	0.7	0.6
Throughput with Min Arrivals (per 15 min)	8.7	12.6	10.7	8.7
Throughput with Max Arrivals (per 15 min)	8.7	7.7	9.5	8.5
Most Common N (aircraft)	8	10	8	21
Max Average N (aircraft)	18	21	12	26

taxi-in time was 6.1 minutes for 31|4 and the maximum was 7.7 minutes for configuration 22|31.

The maximum departure rates for all the configurations range between 0.6 and 0.7 departures per minute, which is equivalent to 36-42 departures per hour. Configurations 22|13 and 4|13 have departure rates of 0.6. The most efficient configuration is 22|31, with a departure rate of 0.7 aircraft/min. It should be noted that the typical number of active departures on the ground seen in this study is significantly less for configuration 22|31 than it is for configurations 22|13 and 31|4. It is therefore unclear if the departure rate would remain greater than those of the other configurations for higher levels of active departures.

The number of departing aircraft on the ground peaks at the same times as the taxi-out times. It is intuitive that as the number of active aircraft increases, so does the taxi time. There are peaks in the number of active departures seen around 0900 hours in all the configurations and between 1600 and 1700 hours in several of them. Figure 4-25 shows the average number of active departures by time of day for each configuration. Configuration 22|31 has consistently lower values of N than the other configurations.

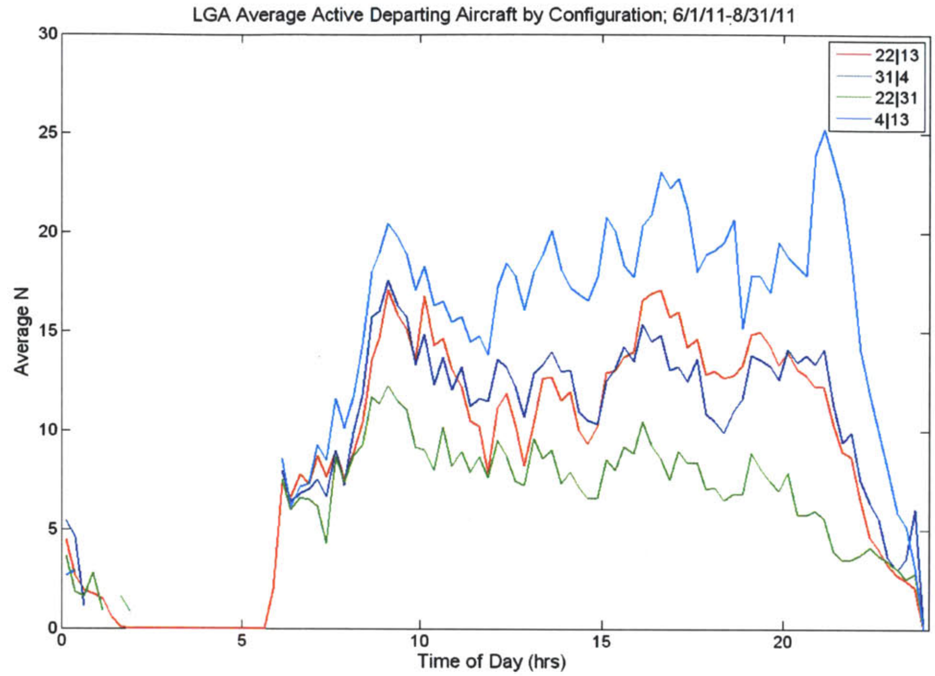

Figure 4-25: Number of active departures by time of day and configuration

The departure throughputs of the configurations with departures on Runway 13 do not vary much based on the number of arrivals. Configuration 22|13 does not vary at all with arrivals and has a maximum rate of 8.7. The departure rate for configuration 4|13 only decreases from 8.7 to 8.5 as the number of arrivals increases. It is expected that these configurations would not vary much, because the arrival runway does not interfere with departures. For 22|13, arrivals on Runway 22 can pass through the intersection within the normal inter-departure separation time. In configuration 4|13, arrivals on Runway 4 will land and hold short of the runway intersection. By contrast, configuration 31|4 shows a great deal of variation. The departure throughput decreases from 12.6 to 7.7 as arrivals increase. This is a decrease in almost 5 aircraft per 15 minutes. This is likely because the takeoffs and landings both begin at the far end of the runways, away from the intersection, and require the maximum amount of departure spacing.

For the time period studied, June-August 2011, the dominant configuration was 22|13. However, the data shows that this is not the most efficient configuration, but is probably warranted by the summer weather conditions. It is an efficient configuration for departures because of the minimum conflict with arrivals. In addition, there is sufficient room for queuing for Runway 13.

4.8 Conclusion

After discussing an aggregate analysis of LGA in the previous chapter, this chapter performed an analysis for each of the four main runway configurations, namely 22|13, 31|4, 22|31 and 4|13. The queuing behavior for each was discussed, based on the manual observation of the airport visualizations under that configuration. Then the departure and arrival taxi times were discussed, noting the maximum and average values throughout the day. Finally the departure rates and throughput were discussed. Each of those was calculated based on N, the number of active departures on the ground. The throughput was found to vary quite a bit based on the number of arrivals, as seen for configuration 31|4.

This process for data manipulation and analysis of ASDE-X data can now be repeated for other airports for comparison purposes. The next two chapters will present an analysis of Philadelphia International Airport.

Chapter 5

Analysis of Operations at PHL

5.1 Introduction

Philadelphia International Airport (PHL) is the only major airport serving the greater Philadelphia area, which is the 5th largest metropolitan area in the United States. In 2011, PHL accommodated 30.8 million passengers, including 4.3 million international passengers, and handled 448,129 aircraft takeoffs and landings [8]. There are 29 airlines which have operations out of PHL, although the predominant carrier by far is USAirways with over 60% of the operations there. Figure 5-1 shows the airport diagram for PHL including runways and terminals. It can be seen that PHL has four runways. Runways 9R/27L and 9L/27R are parallel runways, and are the longest of the runways at 10,500 ft and 9,500 ft respectively. The majority of takeoffs and landings occur on these two runways. In general, when the airport is operating in a west flow, takeoffs occur on 27L and landings on 27R. In an eastbound configuration, 9L is used for departures and 9R for arrivals. Runway 17/35 intersects with 9L/27R and is 6,500 ft long. It can be used for arrivals or departures. Runway 8/26 is the shortest runway at 5000 ft, and is mostly used for general aviation aircraft. The airport has seven terminals: A-West, A-East, B, C, D, E and F. In addition, UPS has a large cargo operation on the opposite side of the runways as the main passenger terminal.

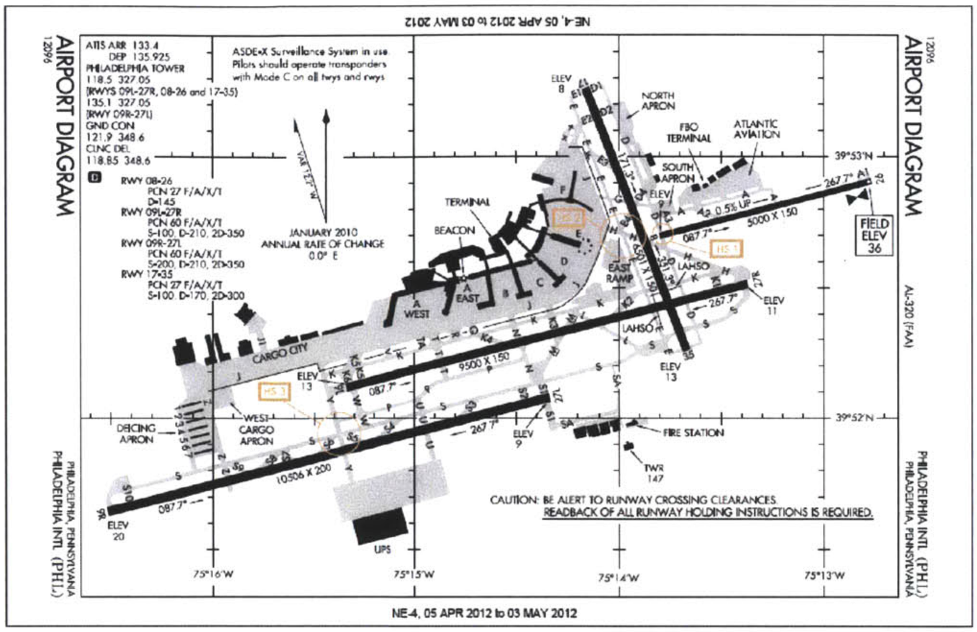

Figure 5-1: Philadelphia International Airport Diagram

5.2 Motivation

Philadelphia International Airport was also selected for analysis because of the amount of surface congestion there. It is known for being an airport with large amounts of delays and long departure queues for the runway. Similar to LGA, the analysis of ASDE-X data for PHL can be used for the characterization of current surface operations, to monitor changes in traffic patterns, and to recommend techniques for surface congestion management. Like LGA, besides just the data analysis, the surface visualizations provide much insight into the patterns of traffic movement.

Another motivation for studying PHL is to compare it with operations at LGA. Both are congested airports, but they differ in a number of ways. First, unlike LGA, PHL is not a slot-controlled airport. This fact would suggest there will be much more variance in metrics like departure counts and taxi times at PHL, because there are no predetermined limits. Another way the airports differ is that PHL has parallel runways which can operate independently of each other, whereas LGA has intersecting runways. This suggests that PHL would have higher capacity thresholds than LGA. In addition, PHL also has two more runways than LGA. These are shorter runways and may not have the same volume of traffic as the parallel runways, but they should further increase the capacity by allowing for more flight operations at the airport.

This chapter discusses operations at PHL overall, and Chapter 6 gives a more detailed description of operations by runway configuration.

5.3 Comparison of Pushback Times

The same analysis of pushback times was completed for PHL as for LGA. This was a comparison of ASPM push times versus ASDE-X initial capture times. The ASPM push times are expected to be earlier, since the ASDE-X initial capture often occurs after pushback. The analysis was done for June-August of 2011, first using all the ASPM data, and then using only the OOOI flights.

As before, individual flight pushback times were compared in ASPM and ASDE-X, by comparing callsigns of departing aircraft for each day in each data source. For PHL, there were 89% matching flights between the two datasets. The two pushback times of each matching flight were recorded, and the differences were calculated. The results can be seen in a histogram in Figure 5-2, grouped into bins of 5 minute intervals. A negative value signifies that the ASPM data had an earlier recorded time.

It was also recorded for each flight whether the ASPM time was automatically recorded (OOOI = Y) or not. The same analysis was performed only using these flights. OOOI flights are roughly 43% of all flights. Figure 5-3 shows the distribution of time differences for all OOOI flights grouped into bins of 5 minutes each.

The average difference in pushback times for all flights was -5.3 minutes, while that of the OOOI flights was -6.5 minutes. The ASPM push time was earlier in both cases, as expected.

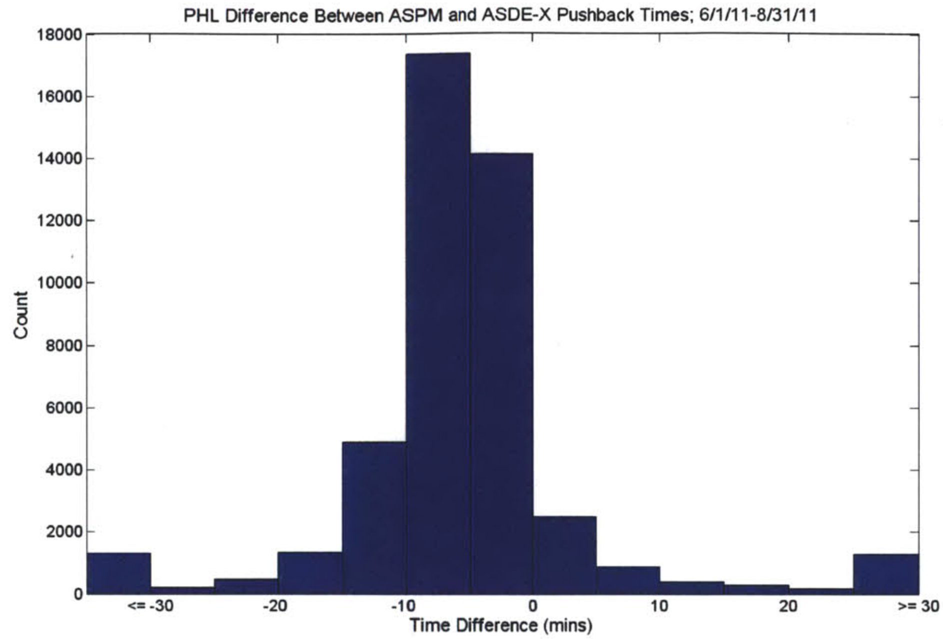

Figure 5-2: Difference between ASPM and ASDE-X push times; PHL Summer 2011

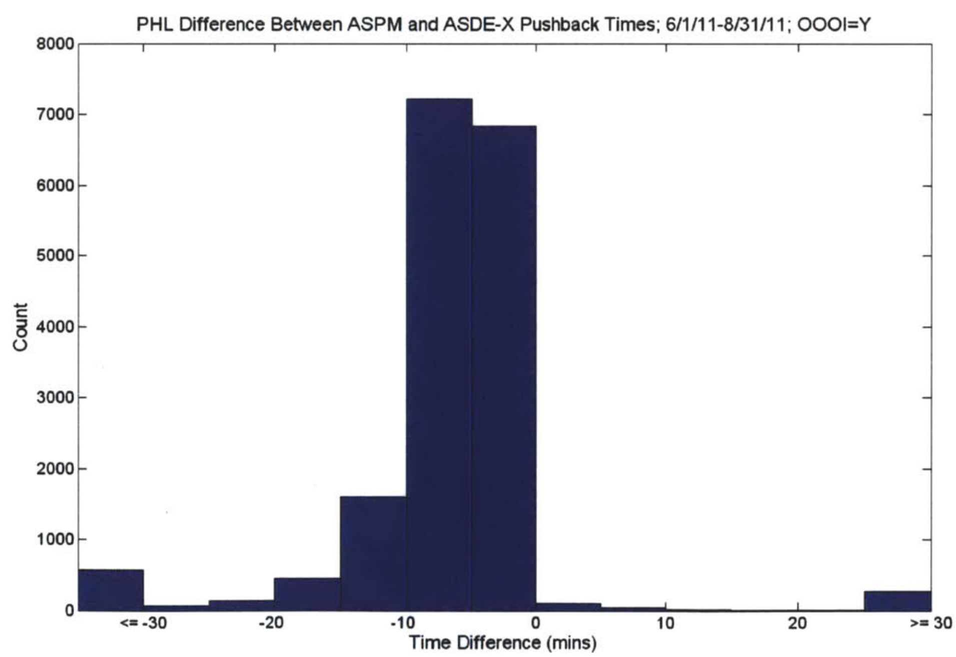

Figure 5-3: Difference between ASPM and ASDE-X push times when OOOI=Y; PHL Summer 2011

5.4 Configuration Usage

As stated above, PHL generally operates with the majority of flights on the two parallel runways, either eastbound or westbound. Typically, one of the runways will be used solely for departures, and one solely for arrivals. When operations are eastbound, Runway 9L is used for departures and Runway 9R is used for arrivals. In addition, Runway 17/35 may be used for arrivals and departures. Runway 8 may also be used for departures. When operations are westbound, Runway 27L is used for departures and Runway 27R is used for arrivals. Runway 17/35 may also be used for arrivals and departures. Runway 26 is used for arrivals. No departures take place on Runway 26 or arrivals on Runway 8, because the terminal building interferes with the flight path. Figure 5-4 below shows the percentage of time that each configuration is in use.

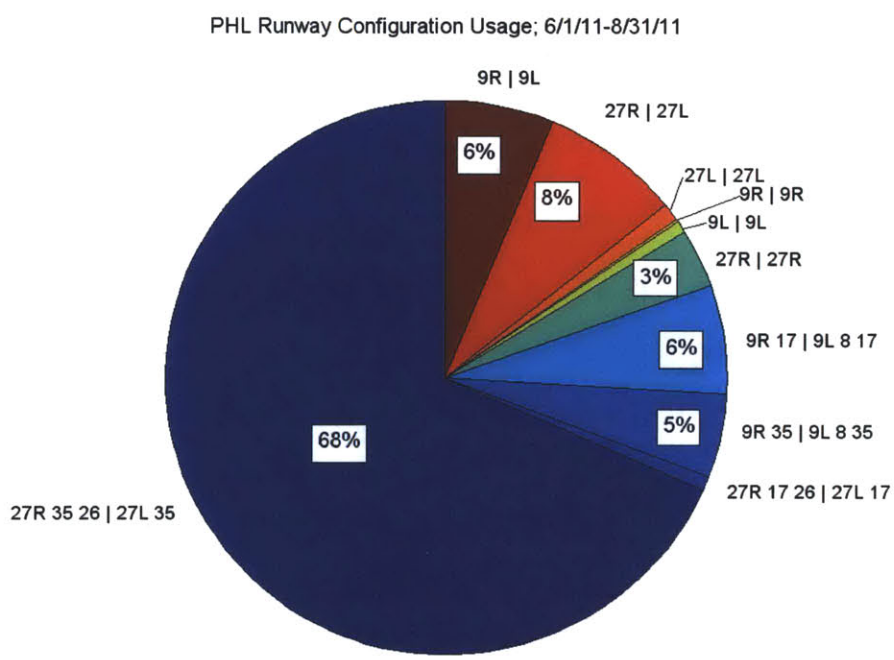

Figure 5-4: PHL runway configuration usage

In total, operations occur with arrivals on 27R and departures on 27L about 77% of the time. Within that, 68% of the time there are also arrivals and departures on Runway 35, in configuration 27R,35,26 | 27L,35. Only 1% of operations occur with arrivals and departures on Runway 17 in configuration 27R,17,26 | 27L,17. The remaining 8% is when operations are only on 27R | 27L, with no operations occurring on the other two runways.

Operations occur in an eastbound flow, with arrivals on 9R and departures on 9L, 17% of the time. In this case, there is a fairly even split between the usage of Runway 17 versus Runway 35. Runway 17 is used in conjunction with operations on 9L and 9R 6% of the time in configuration 9R,17 | 9L,17,8. Runway 35 is used only slightly less at 5% of the time in configuration 9R 35 | 9L 17 8. Finally, the parallel runways 9R and 9L are used alone, without any operations on the other two runways, a total of 6% of the time.

One other runway configuration of note is configuration 27R | 27R, when mixed operations occur on Runway 27R. It is used a total of 3% of the time, and has been observed to occur only between 1000 and 1200 hours on some weekdays, for about an hour. An example of this configuration change from 27R | 27L to 27R | 27R and back is shown in Figure 5-5 below. The box in red on each graph is showing the time period where a change occurs from departures on 27L to 27R. This tends to be a time when both arrival and departure demand are low, so it is easy to accommodate mixed traffic on one runway. It is possible that the airport uses this time to perform runway maintenance on 27L.

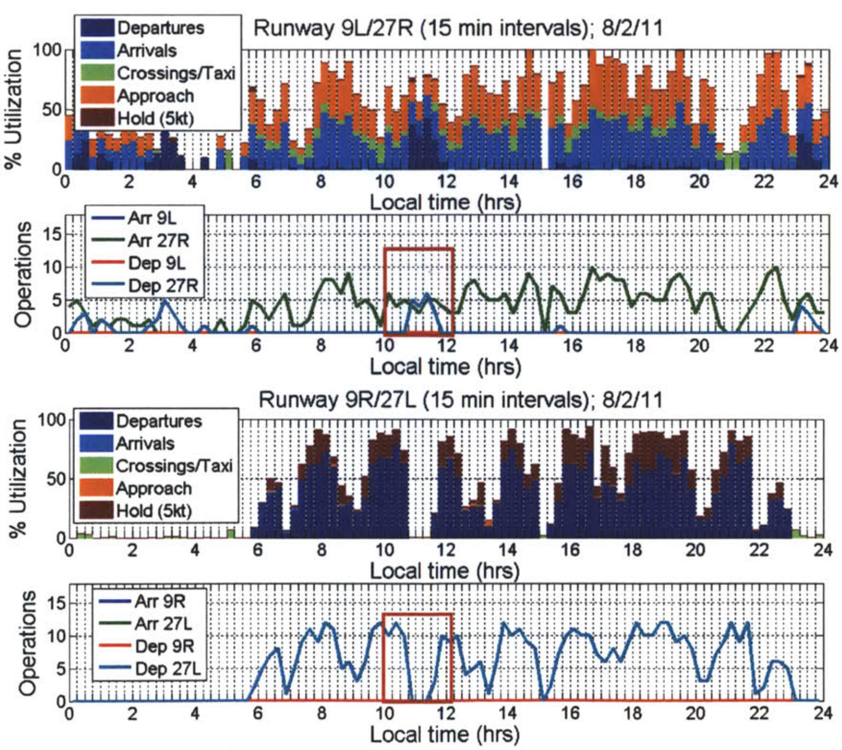

Figure 5-5: Configuration change from 27R | 27L to 27R | 27R

5.5 Queue Boxes

As was done for LGA, queue boxes were determined for the runways at PHL to determine how many aircraft are waiting in queue for the runway threshold at any given time. The queue boxes are also necessary to determine how long aircraft wait in queue, and what percentage of total taxi-out time is spent waiting for the runway threshold.

Most departures occur either on 9L or 27L, so these are the runway departure queues that were studied most closely. However, queues boxes were defined for all remaining runways. The only runway where departures never occur is Runway 26 because the flight path is blocked by the main terminal building. Departures on 9R are rare, but have been seen to occur during mixed-use operations on 9R.

Once again, the queue boxes were determined by watching airport visualizations and seeing where the queues formed for each runway. These were observed both in low demand and high demand states. More attention was given to the time periods with high demand, both because these were the times with the longest queues, and also because it was useful to see how the congestion affected taxi behavior, both for departures and arrivals. The periods of low to medium demand tend to have more consistent queuing and taxi patterns, albeit different than for high-demand periods. During low demand periods, aircraft will typically take the shortest path to the departure runway. However, in high demand periods, it is often necessary to send aircraft on a longer taxi path so they can queue for the runway in places where there is more room so that the queued aircraft will not be in the way of either arriving aircraft or gate operations.

Departures on 27L require crossing the arrival runway 27R, so queues are sometimes observed waiting to cross the runway. Queue boxes were also defined for these runway crossing queues. In the opposite direction, there are no runway crossing queues because departures take place on Runway 9L, which is closest to the terminal. The arrivals on 9R will have to cross 9L to taxi back to the ramp area.

As was the case with the LGA analysis, overflow queue boxes are defined for some runways at PHL. The overflow queue boxes only become active when the queue box immediately prior to it becomes full. What classifies as being "full" is different for each queue box, and was defined by manually observing operations through the visualizations to see when aircraft start queuing in the overflow queue box. Again, the reason for activating the second queue box only when the first is full is to prevent unnecessarily including aircraft which are not waiting in the queue, such as aircraft being held on the ground for non-congestion related reasons.

5.5.1 Runway 27L

As shown above in the runway configuration usage plot, Runway 27L is used for departures 77% of the time. It is, by far, the runway with the most departures, and also the one with the most complex queuing behavior. There are five taxiways feeding into the Runway 27L threshold, with none seeming to be the primary runway queue. Figure 5-6 below shows a visualization of the queue boxes for Runway 27L. The main queue boxes are shown in green while the overflow queue boxes are drawn in blue.

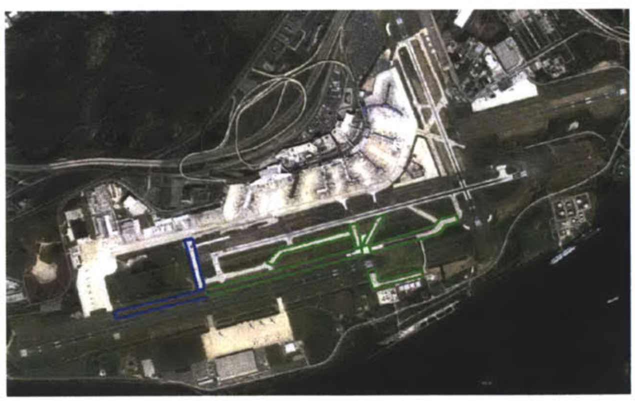

Figure 5-6: Queue boxes for Runway 27L

5.5.2 Runway 9L

Figure 5-7 below displays a visualization of the main and overflow queue boxes for Runway 9L. There are only two taxiways that feed directly into the runway threshold. However, there can be two queues running parallel to each other, and to the runway on the terminal side of Runway 9L. As these get congested, it affects the movement of aircraft in the ramp areas. Therefore, aircraft will be sent to the west, down and around the end of Runway 9L to approach it from the other side. This is a longer taxi route for the aircraft, but it will alleviate congestion in the ramp area.

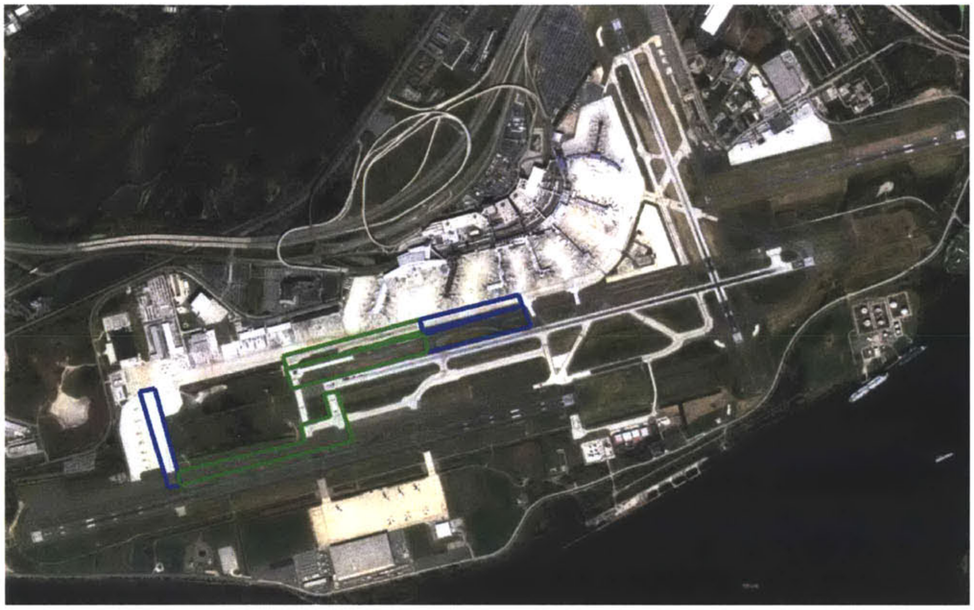

Figure 5-7: Queue boxes for Runway 9L

5.5.3 Runways 17, 35, 8, 9R and 27R

The remaining queue boxes are for the runways which are not the primary departure runways. Runways 17, 35, and 8 do not exhibit long queues, and rarely have more than five aircraft in queue. The queue for Runway 35 is not at the end of the runway, but rather on the opposite side of the intersection with Runway 9L/27R. Most departures on Runway 35 are smaller aircraft which do not need the full runway length. This avoids any interference with the arrivals or departures occurring on 9L/27R. Runways 9R and 27R have queue boxes as well, for the times when they operate as mixed-use runways. Figure 5-8 presents a visualization of all the remaining runway queues.

Figure 5-8: Queue boxes for Runways 17, 35, 8, 9R and 27R

5.6 Propeller-driven Aircraft

As described in Chapter 3, propeller-driven aircraft can affect operations at an airport, since they may be handled differently from jet operations, with different spacing requirements. Therefore, a significant amount of propeller-driven aircraft can affect factors such as throughput, because it may be possible to insert a prop in between two jets, with little to no time penalty for spacing requirements.

As in the case of LGA, PHL does not have a significant number of propeller-driven aircraft operations. In the three months studied, there were over 115,000 flights of which only 589 were props, equivalent to 0.5% of all operations. The average number of props per day in this time period was 6.4. On any given day, the maximum number seen was 17 prop departures and arrivals. These numbers are not large enough to suggest that there would be any change in the departure rate due to the effect of propeller-driven aircraft.

5.7 Overall Analysis

5.7.1 Taxi Times

This section examines the average departure taxi times for PHL over the course of a day. The taxi-out times are binned into 15-minute time intervals, and averaged across all flights beginning to taxi in that interval. The overall analysis includes all configurations, but will be separated by configuration in the next chapter. Figure 5-9 shows two graphical representations of the taxi out times.

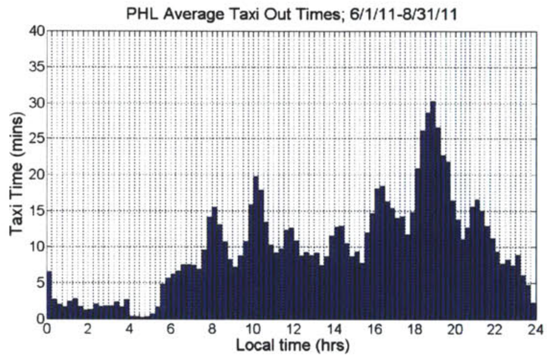

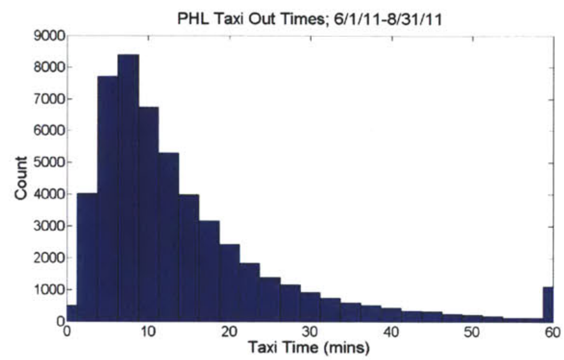

Figure 5-9: (Left) PHL average taxi-out times across all configurations; (Right) PHL distribution of taxi times

The average departure taxi times from all configurations, broken down into 15-minute time periods, are shown on the left in Figure 5-9. Similar to the arrival and departure count plots, this graph also exhibits several small peaks throughout the day. The maximum in the morning is at 1030 hours with an average taxi time of 20 minutes. In the evening there is a large peak in the taxi out times. The maximum for the day occurs at 1900 hours with an average time of 30 minutes. This corresponds to the time period shown above during which the departure and arrival banks overlap. The overall volume of traffic results in high taxi-out times in the early evening. The taxi-out times drop steadily after 2100 until the overnight hours.

The righthand graph of Figure 5-9 shows the breakdown of all the taxi times into bins of 2.5 minutes each. The most frequently observed departure taxi times are around 7.5 minutes long, with over 8,000 occurrences over the three months being studied. The calculated average departure taxi time of all flights in the study is 15.2 minutes.

The arrival taxi times were also studied to see if the congestion on the ground affects the length of the arrival taxi times. Figure 5-10 shows the arrival taxi times for all configurations at PHL. Like with LGA, there is little variation in taxi-in times throughout the day, and the average for all flights is 4.2 minutes.

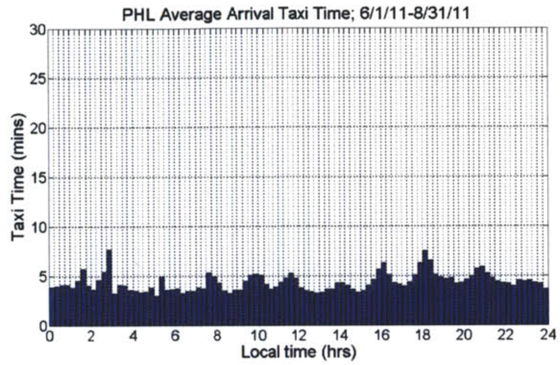

Figure 5-10: Average arrival taxi times across all configurations; PHL Summer 2011

5.7.2 Arrival and Departure Counts

The average arrival and departure counts for PHL were calculated to observe any variations throughout the day. This was an aggregate analysis including all runway configurations. The average arrival and departure counts per 15-minute time period are plotted in Figures 5-11 and 5-12. These are indicative of the traffic levels throughout the day.

Unlike LGA, which has relatively constant levels of arrivals and departures throughout the day, traffic levels fluctuate at PHL with several peaks and valleys. The average departure counts can be see in Figure 5-11. The departures begin at 0600 hours and quickly taper off after 2200 hours. There are peaks and low points of demand throughout the day, with seven peaks in traffic being seen. Each peak is at least 8 departures per 15 minutes, with the maximum being an average of 11 departures per 15 minutes and occurring at 1000 hours. There is a sustained level of over 6 departures between 1600 and 2000 hours, showing sustained traffic in this afternoon to evening time period.

The average arrival behavior over all the runway configurations is shown in Figure 5-12. The arrivals ramp up shortly after 0600 hours and dwindle down after 2200 hours. Arrivals can be seen until 0200. Like the departures, the arrival counts show peaks and valleys in the

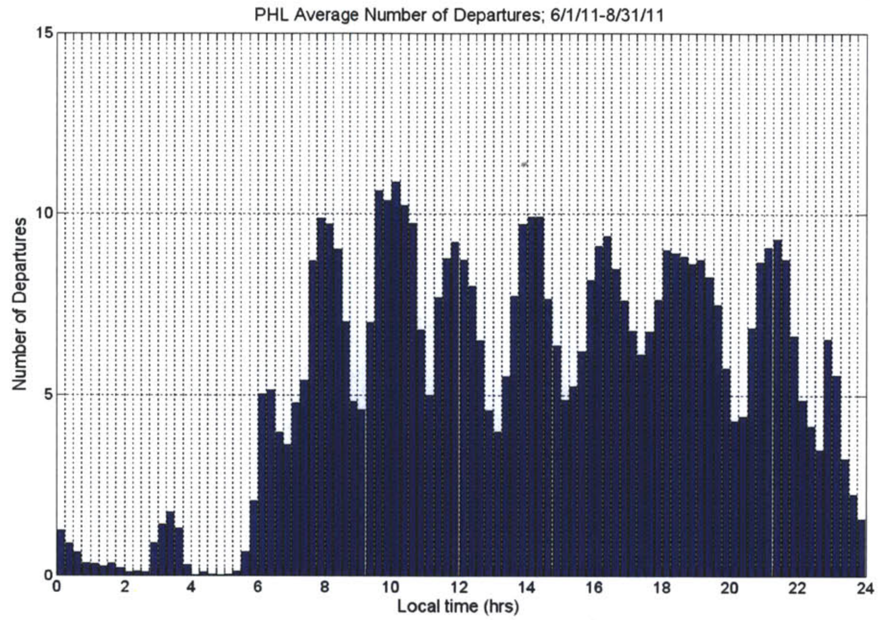

Figure 5-11: Average number of departures across all configurations; PHL Summer 2011

demand rate. The maximum average arrival rate is 13 aircraft per 15 minutes. The average arrival rate stays above 7 per 15 minutes between 1230 and 1730 hours.

Figure 5-13 shows the average departure and arrival counts overlaid on top of each other, showing the relative timeing of the arrival and departure banks. They are generally offset from each other, with the arrivals arriving an hour before the departures depart. It is only around 1600-1800 hours that the two overlap, which is also when the most traffic on the ground is seen. It makes sense that this is the time when the taxi times are highest, because this is when departures and arrivals have the most interaction with each other. Lower departure rates may be expected in this time period because of the number of arrivals. However, given that the two runways are operating independent of each other, there is little effect, as will be shown in the throughput plots in the next chapter.

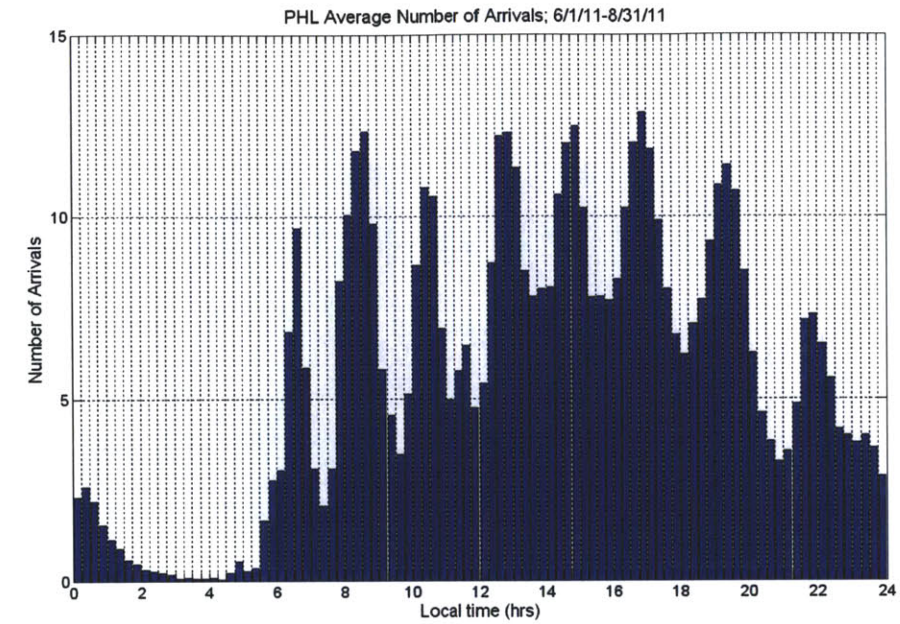

Figure 5-12: Average number of arrivals across all configurations; PHL Summer 2011

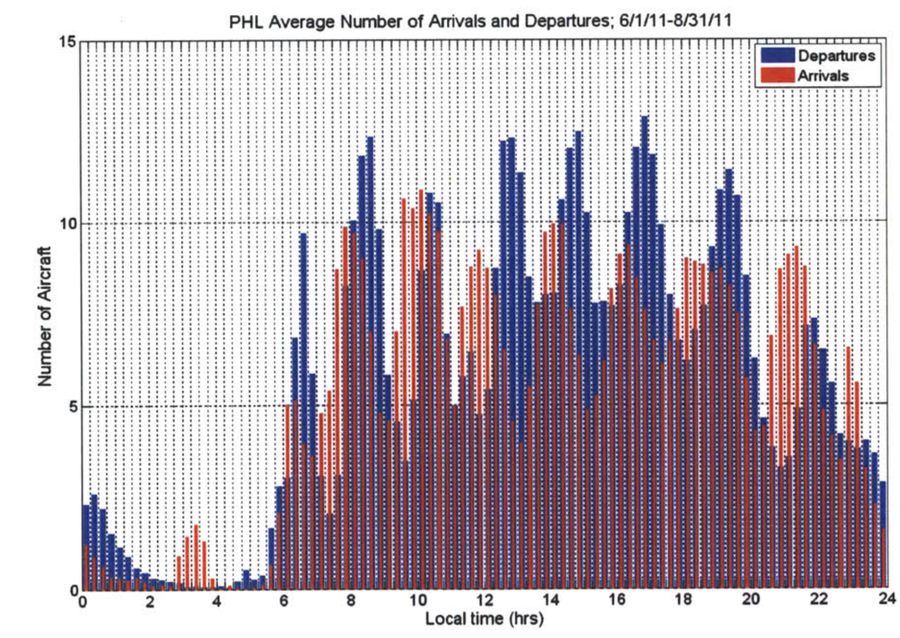

Figure 5-13: PHL arrival and departure banks

5.8 Comparison of PHL and LGA Airports

After completing an analysis of the aggregate data for PHL, we compare PHL and LGA and discuss the reasons for the differences.

To begin with queuing behavior, LGA has one primary long straight queue for three of the four runways. Runway 13 does have several taxiways in the main queue box, but they all feed into one final queue, so only one approaches the runway threshold. The main departure runway at PHL, Runway 27L, has five taxiways that approach the runway threshold. This gives the local controller more flexibility for sequencing departures. Also if the aircraft that is first in line for the runway cannot depart for some reason, it is easy to fill that slot with another aircraft quickly. At LGA that kind of situation will likely cause a decrease in runway throughput, because the runway will be underutilized for a small period of time as the aircraft are resequenced.

The arrival and departure count graphs for LGA and PHL also look quite different. PHL has several banks of aircraft arrivals or departures, with sharp dips in demand between successive banks. LGA demand tends to be more constant throughout the day, with no significant peaks. The difference in the peaking behavior is likely due to the implementation of slot-control restrictions at LGA, which causes the airline schedules to be capped to a maximum capacity each hour, and flights to be spread evenly throughout the day. PHL has no such restrictions, and therefore departures and arrivals occur in banks. The maximum runway utilization at LGA and PHL are 9 arrivals/9 departures and 13 arrivals/10 departures respectively.

Both airports exhibit peaks in taxi times in the afternoon, with the longest taxi times being seen around 1600 hours at LGA and around 1900 hours at PHL. The maximum average taxi times are similar, with only two minutes difference between the two airports. As with the departure and arrival counts, PHL has peaks and valleys in taxi-out times, as departure demand increases and decreases throughout the day. Taxi out times at LGA exhibit a peak in the morning and afternoon, but stay high throughout the rest of the day. At PHL, taxi-out times will drop below 10 minutes during off-peak times. Taxi-out times do not drop below 15 minutes for the bulk of the day at LGA. PHL has a shorter average taxi-out time across all flights of 15 minutes, compared to 20 minutes at LGA. At both airports, the arrival taxi times are short throughout the day.

Table 5.1 provides a summary comparison of several of the calculated metrics discussed in this chapter.

Table 5.1: Comparison of LGA and PHL Calculated Metrics

	LGA	PHL
Pushback Time Difference–All (min)	6.4	5.3
Pushback Time Difference–OOOI (min)	7.7	6.5
Average Taxi-out Time (min)	20.3	15.2
Max Average Taxi-out Time (min)	28	30
Average Arrival Time (min)	6.8	4.2
Max Average Departures (per 15-min)	8	11
Max Average Arrivals (per 15-min)	9	13
Avg Props (per day)	2.4	6.4
Max Props (per day)	8	17

5.9 Conclusions

This chapter presented an aggregate analysis of Philadelphia International Airport. The calculations were done using data from all the different runway configurations. The use of ASDE-X data has given some insights into queuing behavior at the different runway thresholds. Runway 27L has five taxiways feeding into the runway threshold. Because it is handling departures on the outer runway from the airport terminal, the queues are able to line up on these taxiways without much impact on other surface movements. The other major departure runway, Runway 9L, only has two taxiways feeding into the runway threshold. The queue on the near side of the runway has the shortest taxi time, but can back up into the ramp area and impact the movements of both arrivals and other departures.

The average number of departures and arrivals both have several peaks throughout the day, with departures maxing out at an average of 11 per 15 minutes, and arrivals reaching a maximum average rate of 13 per 15 minutes. For most of the day, these departure and arrival banks are staggered. They occur at the same time in the afternoon between 1600 and 1800 hours, which is when the most traffic is seen. This results in long queues and increased taxi-out times for this period of time in the late afternoon. The changes in departure demand do not appear to interfere with arrivals, and arrival taxi times stay constant throughout the day.

Lastly, PHL operations were compared with those at LGA. Some reasons for the differences between the two are the different runway layouts, with PHL having parallel runways and LGA having intersecting ones, and the use of slot controls at LGA.

In the next chapter, PHL configurations using both of the main departure runways will be studied, to see how the metrics differ from one configuration to another.

Chapter 6

Analysis of PHL Configurations

6.1 Introduction

As stated in the previous chapter, Philadelphia International Airport operates either in an eastbound or westbound configuration. When operating westbound, the arrivals are predominantly on 27R, and departures are on 27L. There may also be arrivals and departures on Runway 17/35, more commonly Runway 35. In addition, there may be arrivals on Runway 26. The airport was in one of these configurations 77% of the time during the time period studied, June-August 2011. There are also times when the airport is operating westbound, with mixed-operations on Runway 27R, and Runway 27L not being used. This occurs 3% of the time.

For eastbound operations, the typical setup has 9R as the main arrival runway and 9L as the main departure runway. Again, there may be departures and arrivals on Runway 17/35, being used roughly the same amount in each direction. There may also be departures on Runway 8. The airport was in one of these configurations 17% of the time. Besides this usual arrangement, there are times of eastbound operations when either Runway 9R or Runway 9L is being used for mixed operations.

Rather than studying each specific runway configuration, this chapter characterizes airport operations collectively for eastbound and westbound configurations. Eastbound configurations include any configuration in which arrivals are on 9R and departures are on 9L. Westbound configurations include any configuration in which arrivals are on 27R and departures are on 27L.

Together, these configurations are used 94% of the time. By using ASDE-X data to study these two sets of configurations, the vast majority of operations at the airport will be categorized and analyzed.

6.2 Configuration 27R|27L

The most commonly used runway configurations have arrivals on Runway 27R and departures on Runway 27L. The sections below will discuss queuing behavior, taxiing, arrivals, departures and throughput for this group of configurations.

6.2.1 Queuing Behavior

The queuing behavior for Runway 27L is not as simple as that of the LGA runways. There are five separate feeder taxiways leading to the runway threshold. Where aircraft line up is situational, with no apparent consistent procedures on which aircraft queue where. There are two taxiways running parallel to the runway, one east and one west of the runway threshold. These two taxiways tend to have the longest queues, with aircraft being placed on one side or the other based on their starting location at the gate. Gates on the eastern side of the airport will go down the taxiway parallel to Runway 17/35, cross Runway 27R and then approach the runway threshold from the east. When there is a significant number of departing aircraft, there can be a runway crossing queue waiting to cross Runway 27R. Gates on the western side of the airport will go all the way around Runway 27R and approach from the west. The taxiway feeding in from the south is occupied either by small aircraft like regional jets, or by flights with short holds [9]. The small aircraft can sometimes fit into an opening that other larger jets would not be able to. The short holds are aircraft with departure times in the near future. Figure 6-1 shows an example of the departure queues for Runway 27L.

Figure 6-1: Queue for configuration 27R|27L

6.2.2 Taxi Times

The taxi-out times for the 27R|27L configurations appear similar to the overall taxi-out times. There are several small peaks throughout the day, corresponding to the multiple departure banks, with the largest peak occurring around 1900 hours. The maximum average time in the morning is attained around 1015 hours, with a taxi-out time of 17 minutes. The longest average departure taxi time is at 1900 hours with a length of 28 minutes. The overall average taxi time of all aircraft is 14.3 minutes. The graph of departure taxi times by 15-minute interval is shown in Figure 6-2.

The arrival taxi times for configuration 27R|27L do not vary much from the overall arrival taxi times. The average time is 4.0 minutes. The average arrival taxi times for the 27R|27L configurations are shown in Figure 6-3.

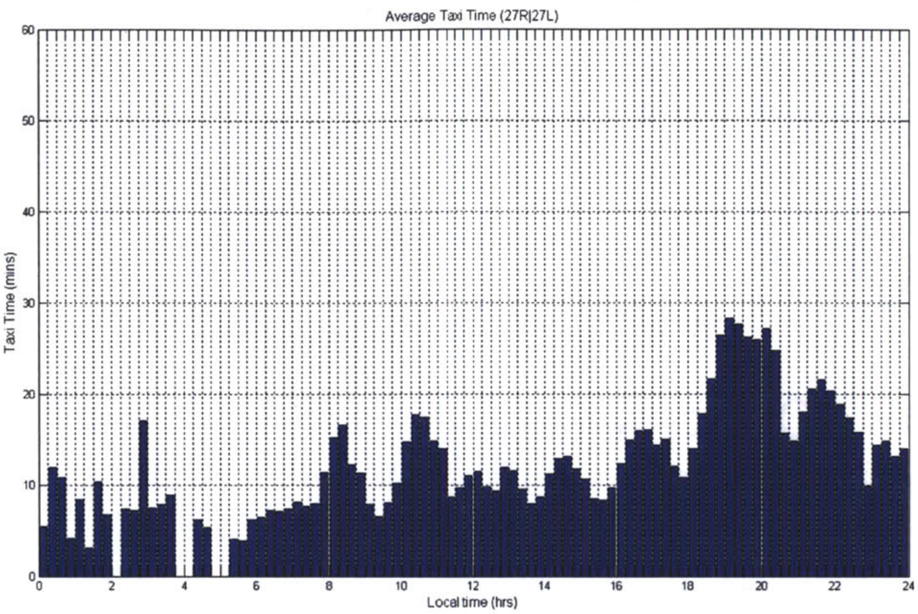

Figure 6-2: Average taxi-out times for configuration 27R|27L

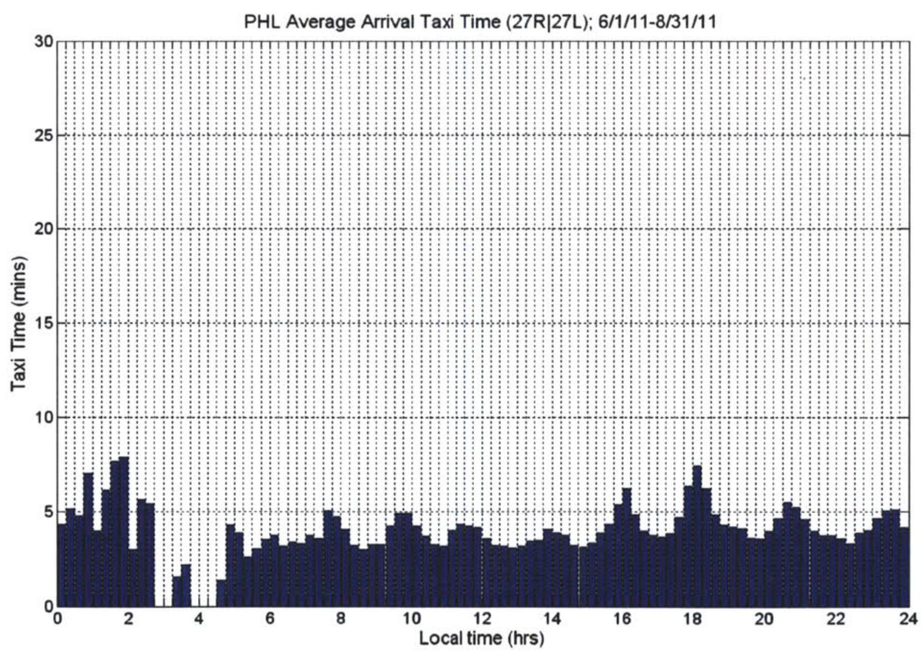

Figure 6-3: Average arrival taxi times for configuration 27R|27L

6.2.3 Departure Rates

As stated in an Chapter 4, the departure rate reflects the number of departures per minute as a function of the number of active departing aircraft on the ground. Figure 6-4 shows the departure rates for all the configurations using 27R|27L. As a reminder, N is the number of departures that have already pushed back and are taxiing or waiting in queue for the departure runway.

For this set of configurations, the maximum departure rate is 0.86, which is equivalent to 52 departures per hour. This occurs when N is equal to 17. The departure rate initially increases, reaching about .8 aircraft per minute when the number of departing aircraft equals 10. Then it stays at or above .8 until the departing aircraft equals 22, at which point it starts decreasing slowly. The most common value of N is 3.

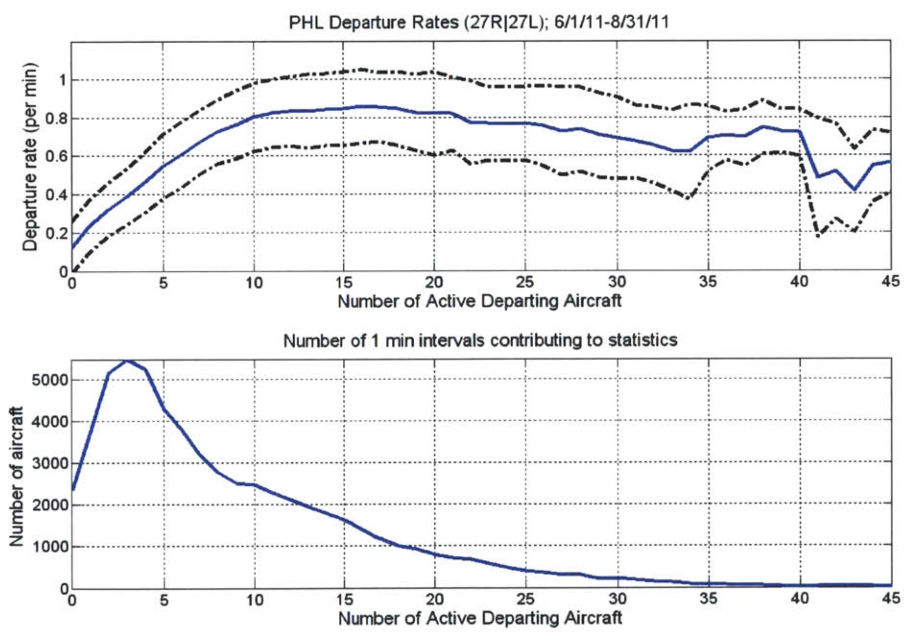

Figure 6-4: Departure rates for configuration 27R|27L

Figure 6-5 shows how the number of active departing aircraft (N) varies by time of day. There are a number of peaks throughout the day, similar to the graph of the number of departures. The number of active departing aircraft is largest around 1900 hours, with an average of 24 aircraft. This time period is also when the taxi-out times are the longest.

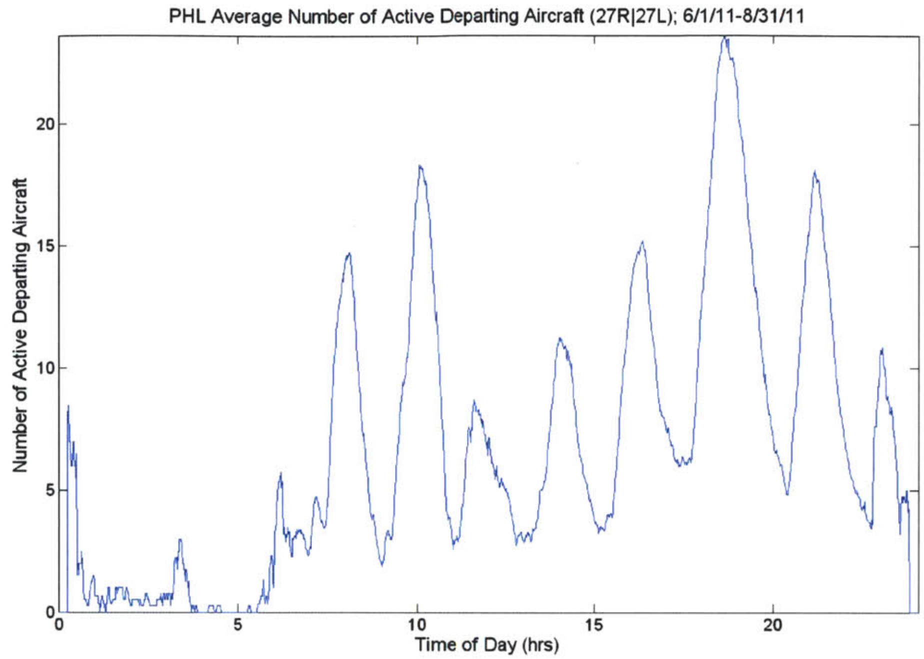

Figure 6-5: Departing aircraft counts for configuration 27R|27L

6.2.4 Throughput

The average throughput parameterized by number of arrivals is shown below in Figure 6-6. In this case, for the set of configurations 27R|27L, the departure throughput does not vary with the number of arrivals. This is because the arrival and departure runways are parallel to each other and far enough apart to operate independently, therefore increasing the number of arrivals should not affect the number of departures. Also, the crossing runway, 17/35, does not intersect with the main departure runway, 27L. The maximum throughput for all levels of arrivals is 12 departures per 15 minutes, or 48 departures per hour. Once the number of active departing aircraft reaches 12, the number of departures holds steady at the maximum value.

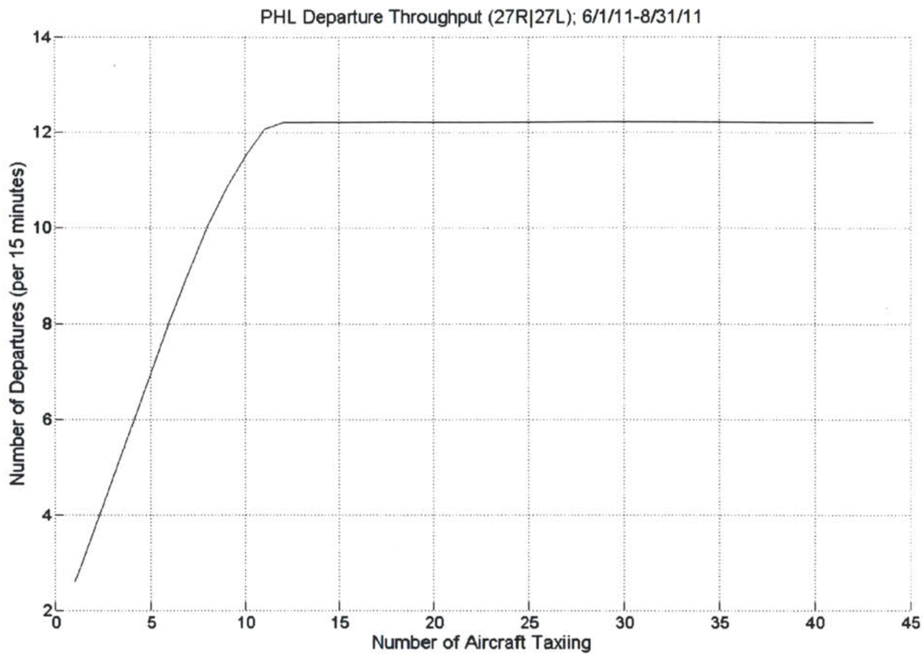

Figure 6-6: Throughput parameterized by arrivals for configuration 27R|27L

6.3 Configuration 9R|9L

The eastbound runway configurations have arrivals on Runway 9R and departures on Runway 9L. This section will discuss the queuing behavior, taxi times, departure rates and throughput for this set of configurations.

6.3.1 Queuing Behavior

For Runway 9L, there are two feeds into the runway threshold, one from the north and one from the south. The feed from the north is actually composed out of queues on two taxiways running parallel to the runway. This is the shortest distance to the runway from the airport gates. However, long queues on both of these taxiways can cause congestion in the ramp area. The queue that is feeding in from the south can help alleviate this congestion. Aircraft taxi down around the end of 9L and then have more space for queuing. This taxi route is longer for the aircraft, but it helps aircraft movements flow more smoothly over the whole airport surface. Arrivals in this configuration typically proceed eastward, cross Runway 9L, and then proceed westbound until they reach their appropriate gate area. This keeps all aircraft moving in a counter-clockwise direction and avoids interactions between arrivals and departures. An example of a queue for Runway 9L is shown in Figure 6-7.

Figure 6-7: Queue for configuration 9R|9L

6.3.2 Taxi Times

The average departure taxi times for the 9R|9L configurations are shown in Figure 6-8. Here, once again, there are several peaks and valleys throughout the day. The taxi-out times are higher overall compared to the 27R|27L set of configurations, and the peaks are more pronounced. There are two morning peaks, with the maximum occurring at 1045 hours, with an average taxi-out time of 34 minutes. The evening peak here occurs a little later than the overall evening peak, with the maximum average taxi-out time being 42 minutes and occurring at 1945 hours. The average taxi time across all departures in this configuration is 22.3 minutes.

The arrival taxi times are shown in Figure 6-9. The average taxi time for all arrivals under this configuration was 5.5 minutes.

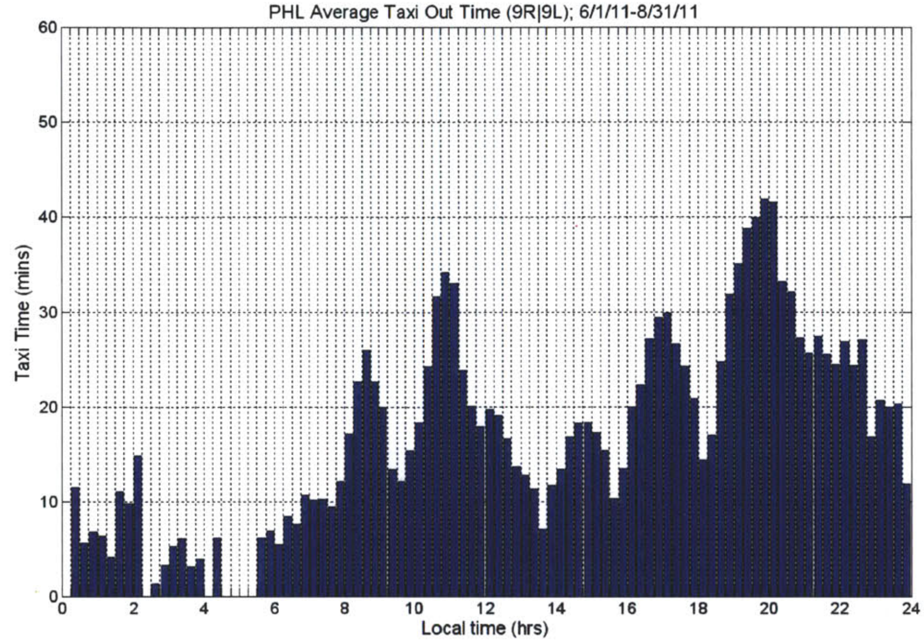

Figure 6-8: Average taxi-out times for configuration 9R|9L

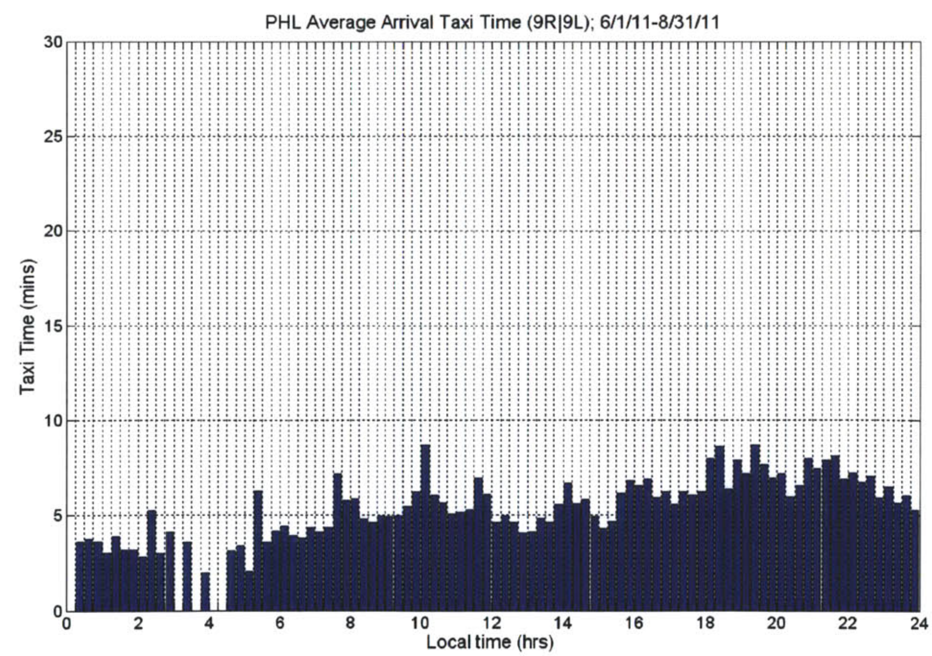

Figure 6-9: Average arrival taxi times for configuration 9R|9L

6.3.3 Departure Rates

The departure rates for the 9R|9L configurations are shown in Figure 6-10. The rate increases steadily until N reaches 15, with a rate just over 0.65 departures per minute, or 39 departures per hour. Beyond that the rate continues gradually increasing until N reaches 31 and a rate of 0.74 departures per minute, or 44 departures per hour. With values of N beyond 33, the departure rate starts to dip.

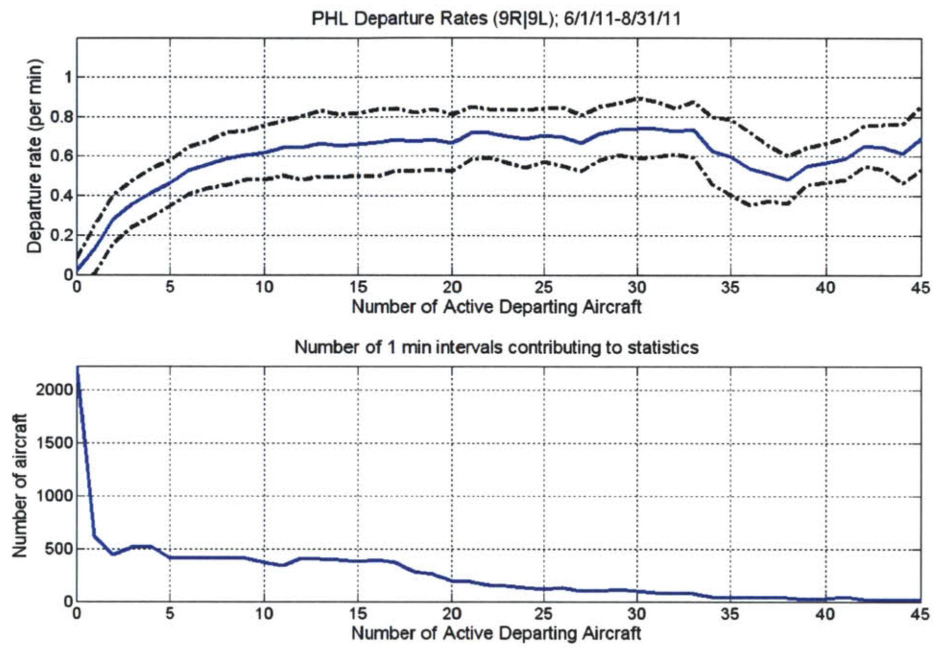

Figure 6-10: Departure rates for configuration 9R|9L

The variation in the number of active departing aircraft throughout the day is shown in Figure 6-11. Once again, there are multiple peaks and valleys in number of aircraft. The shape of the curve closely parallels the rise and fall of departure taxi times. The morning and evening peaks are at 1045 and 1945 hours, respectively. The maximum number of departing aircraft occurs at 1945 hours, with an average of 26 active aircraft.

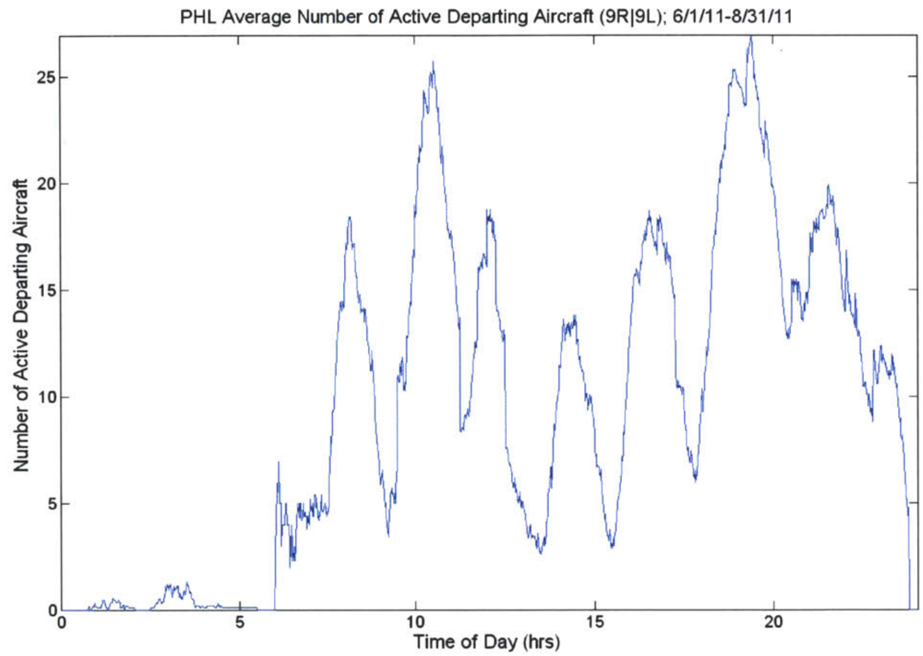

Figure 6-11: Departing aircraft counts for configuration 9R|9L

6.3.4 Throughput

The throughput for configuration 9R|9L also does not vary significantly by the number of arrivals. This is likely because the main arrival and departure runways are parallel and spaced far enough apart to operate independently. The throughput reaches a maximum rate of 10.2 departures per 15 minutes as *N* increases and holds steady at that rate. This is 2 aircraft fewer per minutes than in configuration 27R|27L.

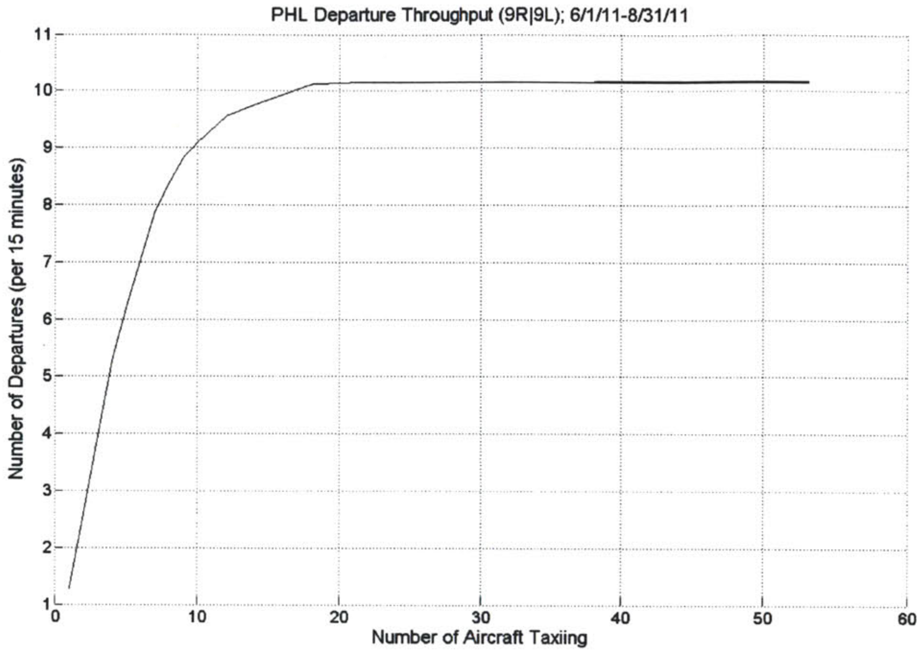

Figure 6-12: Throughput parameterized by arrivals for configuration 9R|9L

6.4 Analysis

A summary of the statistics for each configuration is shown in Table 6.1.

Table 6.1: Comparison of PHL Runway Configurations

	27R|27L	**9R|9L**
Usage	77%	17%
Average Taxi-out Time (min)	14.3	22.3
Max 15-min Average Taxi-out Time (min)	28	34
Average Arrival Taxi Time (min)	4.0	5.5
Saturation Point (aircraft)	22	33
Max Departure Rate (per min)	0.86	0.74
Throughput with Min Arrivals (per 15 min)	12.2	10.2
Throughput with Max Arrivals (per 15 min)	12.2	10.2
Most Common N (aircraft)	3	4
Max Average N (aircraft)	24	27

Runway 27L has five feeder taxiways, whereas Runway 9L only has two. The queuing behavior for both is primarily situational, with no definitive procedures. The main queues for Runway 27L are to the west and east of the runway threshold, and are fed by aircraft coming from gates on the west and east sides of the airport. The main queue for 9L begins on the taxiway parallel to the runway, just north of it, because this is the shortest route to the runway from the gates. However, once this queue grows in size, it begins to congest the airport ramp area. Then aircraft are directed down and around to the other side of the runway. Neither runway has specific holding locations for EDCTs and other held aircraft. The location is situational, depending on where there are empty spaces on the airport surface. Air traffic controllers prefer the 27L arrangement, because there are more aircraft to choose from for the next departure. This allows for an easier job in alternating between north- and southbound departure fixes.

Each configuration shows several peaks in departure taxi times throughout the day. Some are more pronounced than others. There are morning and evening peaks in taxi times around 1000 and 1900 hours respectively. The configuration with the lowest average taxi-out time was 27R|27L with a length of 14.3 minutes. This is 8 minutes less on average than for 9R|9L. The maximum average taxi-out time was 34 minutes for configuration 9R|9L, which is 6 minutes more than for 27R|27L. The average arrival taxi times over a 15-minute time period were fairly constant all day long across all configurations. The average was 4.0 minutes for 27R|27L, and 5.5 minutes for configuration 9R|9L.

The maximum departure rates ranged between 0.74 departures per minute for 9R|9L and 0.86 departures per minute for 27R|27L, which is equivalent to 44-52 departures per hour. The departure rate is likely less for 9R|9L because the departure runway in this configuration intersects with Runway 17/35, which will also handle arrivals and some departures. For configuration 27R|27L, there is not interference between the departing runway and this crossing runway. The arrivals on 17/35 necessitate extra spacing between departures on 9L, and lower the overall departure rate.

There are similar peaks in the number of active departing aircraft and the departure taxi times. This is intuitive because more active departures lead to longer queues and therefore longer taxi-out times. There are peaks in the number of active departures seen around 1000 hours in both the configurations, and the maximum number of departing aircraft are seen around 1900 hours in both configurations.

Neither of the configurations shows a change in the throughput levels based on the number of arrivals. This is expected because both are configurations utilizing parallel independent runways. Configuration 27R|27L is the most efficient configuration. It has a throughput of 12.2 aircraft per 15 minutes, which is 2 aircraft more than seen in 9R|9L. The reason it is more efficient is because the departure runway does not intersect with the crossing runway, so operations on the crossing runway do not effect departures. In configuration 9R|9L, the departure runway does intersect the crossing runway, so there is a decrease in the departure rate. Also, all the arrivals must cross the departure runway which may also contribute to the lower throughput levels.

For the time period studied, June-August 2011, the dominant configuration was 27R|27L, which is also the most efficient configuration. This configuration is always used when winds and ceilings allow it [9]. It is an efficient configuration for departures because of the minimum conflict the crossing runway and with arriving aircraft. Also, there is a good deal of space for queuing for Runway 27L so as not to congest the ramp area of the airport.

6.5 Comparison to LGA

In the previous chapter, an overall comparison of the surface movement characteristics of LGA and PHL was presented. That discussion is continued here by looking at additional metrics like departure rates and throughput. The most-used configuration for Summer 2011 will be discussed. Other configurations will also be touched upon. A summary of the statistics from the dominant configuration for each airport is presented in Table 6.2.

As seen in the overall analysis, the taxi-out times for PHL are significantly shorter than for LGA. Here in the dominant configuration taxi-out times are on average 10 minutes shorter in duration, with the average being 14.3 minutes as compared to 24.3 minutes for LGA. Arrival taxi times are also shorter in length at 4.0 minutes, compared with 7.1 minutes at LGA. The arrival taxi times tended to stay constant throughout the day at both airports, while the departure taxi times had peaks in the day. At PHL there were several throughout the day with the main one occurring around 1900 hours. At LGA there was a morning peak around 0900 hours and and afternoon peak around 1800 hours.

The departure rates for PHL are higher than for LGA. This is mainly because it utilizes

Table 6.2: Comparison of Most-Used Runway Configurations

	LGA 22|13	PHL 27R|27L
Usage	37%	77%
Average Taxi-out Time (min)	24.3	14.3
Max 15-min Average Taxi-out Time (min)	30	28
Average Arrival Taxi Time (min)	7.1	4.0
Saturation Point (aircraft)	28	22
Max Departure Rate (per min)	0.6	0.86
Throughput with Min Arrivals (per 15 min)	8.7	12.2
Throughput with Max Arrivals (per 15 min)	8.7	12.2
Most Common N (aircraft)	8	3
Max Average N (aircraft)	18	24

parallel runways, whereas LGA has intersecting runways. Also, depending on the configuration, PHL may have additional departures on one or both of the two shorter runways, which also contribute to an increased departure rate. The configurations shown in the Table 6.2 have departure rates equal to 52 departures per hour for PHL, and 36 departures per hour for LGA. The LGA configuration with the highest departure rate, 22|31, has 42 departures per hour, still 10 less than for PHL.

The departure throughputs for both of these configurations do not vary much based on the number of arrivals. As seen with the departure rates, the throughputs are greater for PHL. For the most-used configurations, the throughput is 49 aircraft per hour at PHL and 35 aircraft per hour at LGA. Only configuration 31|4 can possibly match this rate at PHL, and only when there are no arrivals.

On the whole, because of the parallel runways at PHL and because of the additional two runways, it is able to operate much more efficiently in terms of departure and arrival throughputs.

6.6 Conclusions

After discussing an aggregate analysis of Philadelphia International Airport in the previous chapter, this chapter has performed an analysis for each of the two main sets of runway configurations. The queuing behavior for each was discussed based on the manual observation of the airport visualizations under that configuration. Then the departure and arrival taxi times were discussed, noting the maximum and average values throughout the day. Finally the departure rates and throughput were discussed. Each of those was calculated based on N, the number of active departures on the ground.

Chapter 7

Implementation and Next Steps

7.1 Introduction

7.1.1 Overview

This chapter will provide a framework for future research using ASDE-X data. The objective is to use real-time ASDE-X data to provide decision support in the air traffic control tower, and thereby improve the efficiency of airport operations. In this chapter, a field implementation of a tablet app for surface congestion control will be discussed. This is the first step in the process towards an ASDE-X based decision support tool, and makes a case for the benefits for future work with ASDE-X data.

7.1.2 Motivation

Over the past several years, many research studies and field trials have been devoted to improving the management of aircraft on the airport surface [7][3][5]. Many of these studies have tried to manage the flow of aircraft taxiing on the surface as a means of simultaneously optimizing runway throughput, taxi times and fuel burn on the surface. The developed algorithms have sometimes used data collected through ASDE-X sensors to quantify predicted taxi times and runway throughputs. However, the implementation of such operational initiatives in the control tower often involves manual collection of data and a some effort on the part of controllers. During preliminary studies, researchers may take care of this manual data collection, but to implement the initiative as an everyday operating procedure would increase controller workload.

With all the data being collected in real time by ASDE-X sensors, the vision would be to eventually automate some of the data processing and disseminations. If real-time surface data could be fed into a computer program that could extract the necessary inputs, then an ideal surface management strategy could be calculated without any excess workload on air traffic controllers. The motivation for this tablet app was to demonstrate one step in the process of building such a decision support tool. It was designed to be a tool that controllers could use to perform the pushback rate control strategy on their own, without any additional assistance.

7.2 Field Test of a Surface Congestion Control Strategy

7.2.1 Background

A research study was completed in the summer of 2010 in the air traffic control tower of Boston Logan International Airport [11] to test an algorithm for reducing congestion on the airport surface, by controlling the rate at which planes were released from the gates. The details of the algorithm can be seen here [10]. This study was completed with two researchers: one who was collecting the necessary inputs for the algorithm and one who processed those values through the algorithm to come up with the appropriate pushback release rate. After the proper pushback rate had been calculated, one of the researchers would place a "3x5" color-coded card on the display in front of the ground controller displaying the appropriate rate, for example 3 per 5 mins, signifying that only 3 aircraft were recommended to be released every 5 minutes. Any aircraft beyond that would need to be held at the gate. Figure 7-1 shows the cards and a picture of the location where the cards were placed in front of the controller.

Given the success of this study in Pushback Rate Control in 2010, the next step was to evolve this process so the controllers would (in principle) be able to carry out the process alone without the help of any researchers. This meant creating a program which would

Figure 7-1: BOS rate control setup

generate the desired rate output given the necessary inputs. Some sort of handheld device seemed preferable because the mobility allows flexibility in how, and by whom, the program is used.

After some evaluation, it was decided that the device to be used was the 7.0" Samsung Galaxy Tab. It was large enough to be seen easily, but still small enough to be portable and easily placed near the controller using it. In addition, the Android™ operating system offered an open and accessible application development environment.

7.2.2 Formulation

Given that the necessary inputs to the algorithm were relatively easy to gather, the initial goal was to find any way to automate the process of generating the output rate so the controllers could work on their own. Several alternatives were thought about as means of accomplishing this, with varying degrees of simplicity and utility. The main goal was ease of use. Also needing to be considered are the user interface, the ability to be installed in the control tower and the potential to add features. Following is a list of some possibilities that were explored.

(a) **Lookup table:** A lookup table could come in many forms, but could be something as simple as a spreadsheet. It is the most basic, and easiest to create, with a very simple input-output format. The downside it is not the most user-friendly in terms of display and limits the ability for adding additional features.

(b) **Graphical User Interface:** A GUI could be any sort of computer program that has a functional and appealing user interface. This would be a step over the lookup table, but would require installation on control tower machinery.

(c) **Application software:** The use of application software (an app) offers freedom from interaction with the tower computing machinery. The small size of tablets and other devices which run apps make it a portable solution. This gives the advantage of not tying the software to any one position in the control tower. The touchpad user interface makes it easy to use. An app is also extendable, so more features could be added as needed.

In the end, the app was chosen as the best method for field-testing such an operational strategy. While it required more work to create than a simple lookup table, it also offered the most flexibility. The use of a mobile tablet allowed the app to be used by any controller, and to not be restricted to one location as a larger piece of hardware would be. This is useful because the device might be used at various control positions, depending on what the workload of those controllers is at a given time. It also easily allowed for the addition of new features to the initial simple design.

7.2.3 Inputs

The created Android app calculates the expected takeoff rate and the recommended pushback rate using a lookup table from the pushback rate control algorithm. For this, all the following state variables have to be given as inputs: Runway configuration, weather, expected arrival throughput, jets under ground control, jets under local control, and number or props. These inputs will be discussed in more detail below. It should be noted that since there are a significant number of props operating at BOS, it was needed as an input. The input interface can be seen in Figure 7-2.

Figure 7-2: Android tablet input interface (Rate Control Transmitter)

(a) Runway Configuration: This is a drop-down list of the possible configurations in use at Boston Logan. This only needs to be input when the runway configuration changes.

(b) Weather: This is a drop-down choice of visual meteorological conditions (VMC) or instrument meteorological conditions (IMC). This field only needs to be updated when visibility conditions change.

(c) Arrivals: This is the number of predicted arrivals in the next 15-minute interval. This number is obtained from the Traffic Situational Display (ETMS).

(d) Ground Jets: This is the number of jets under Ground Control at the start of the 15-minute time period. This number can be obtained by counting the number of flight strips stacked up in the rack in front of the Ground Controller. Any props in the flight strip rack should not be counted. Props can be determined by looking at the information on the flight strip.

(e) Local Jets: This is the number of jets under Local Control at the start of the 15-minute time period. This number can be obtained by counting the number of flight strips stacked up in the rack in front of the Local Controller. Again, any props in this flight strip rack should not be counted.

(f) Props: This is the number of props on the ground, either under Ground or Local Control at the start of the 15-minute time period. They are counted separately because they do not have the same departure spacing requirements as jets.

These metrics only need to be updated every 15-minutes. They are all quickly determined by either looking at the flight data screen or by manually counting the number of flight strips seen. Because these numbers are easy to determine, it is not a burden on the controller to update the inputs in the tablet app, and recalculate the rate for the next time period.

The configuration and weather are not subject to frequent changes, and the number of arrivals can be determined from the ETMS display predictions. The number of ground and local aircraft can both be determined by counting the number of flight strips in front of the respective controller and subtracting any props among them. Similarly, the number of props is found by looking at the flight strips in front of both the ground and local controllers.

7.2.4 Outputs

Once a pushback rate has been obtained by running the app, the controller has two choices of how to display that information: The rate control display and the volume control display.

Rate Control Display

In this mode, the output is simply an image of a color-coded pushback rate, showing the number of recommended pushbacks per interval of minutes. Possible rates are: 1 per 3 mins, 1 per 2 mins, 2 per 3 mins, 3 per 5 mins, 4 per 5 mins, 1 per 1 min, and no restrictions. This display simply automated the process performed in 2010, when the researchers showed the controllers physical cards to inform them of the new rate [11]. It is up to the controller to keep track of the time intervals and how many aircraft have pushed back already. If the number of calls for pushback exceeds the rate, an aircraft will be held until the next time interval starts. Again, it is up to the controller to keep track of holding aircraft, and then releasing them when the next time interval begins.

Figure 7-3: Rate control display

Volume Control Display

The volume control display was developed to help the controllers keep track of the number of aircraft that have called and have already been released, and to keep them aware of the time intervals. It was observed during the field trials in 2010 that many controllers used pencil and paper to keep tallies of the number of aircraft released, so as not to exceed the given rate. The volume control display mode attempts to imitate and improve upon what they were already doing on paper. In addition, it gives them visual cues of the passage of time, and the next action required. The volume control display was expected to significantly reduce the workload of the controller running the program, and possibly merge the Boston Gate position with another position.

In the volume control display, the 15-minute time period is broken down into time intervals based on the rate. For example, if the rate is 3 per 5 minutes, then the display will show three lines, each corresponding to a 5-minute time interval, with each line having three aircraft displayed. A time interval becomes active when the current time is within that time interval, and is indicated by a small black arrow to the left of the time interval. Flights can only be released during an active time interval, otherwise pusback positions can only be reserved. If there are unused release slots for a given time interval, those slots will roll over to the next time interval up to a maximum of twice the rate. For example if the rate is 3 per 5 minutes, then the maximum number of slots in any 5-minute time period is six. Any unused slots in excess of this will simply be discarded. These are the actions available in the volume control display (also illustrated in Figure 7-4:

1. Releasing a plane: If a plane calls for pushback, one of the aircraft in that time interval should be selected. It will change in color from black to gray, indicating that it has been released.

2. Reserving a pushback position: If a plane calls for pushback and there are no more slots available in the current time interval, the controller would tell the aircraft to hold, and reserve a slot for it in a future time interval. This is done by selecting that aircraft on the display, whose icon will then be rotated 45 degrees to indicate it has been reserved. When that aircraft is eventually released in the future, the controller clicks on the aircraft image again to indicate that it has been released. The image will then rotate back and change color to gray.

3. Reserving a plane for a future time period: An aircraft slot can be reserved for a future 15-minute time period by clicking on the white space next to that future time period. A rotated aircraft will then appear to indicate a reservation. When the appropriate time period arrives and the new rate has been calculated, that aircraft will appear already reserved. Because the rate for the future time period is not yet known, it will always be broken up into 5-minute time intervals. Once the actual rate is known, the reserved plane will be assigned to the time interval that is closest to the time that it was reserved for.

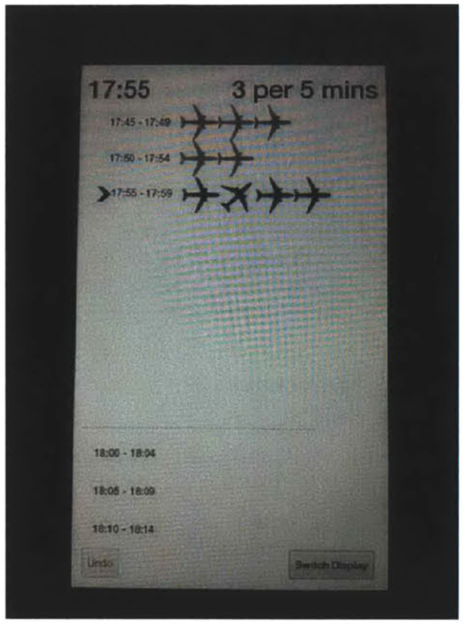

Figure 7-4: Actions in the volume control display

7.3 Deployment and Testing

The given rate is valid for 15 minutes, and then needs to be updated. A timer was included in the app to remind the controllers that it was time to update the inputs and rerun the control algorithm to receive a new rate. Depending on the workload, it may be too difficult for the controller who is releasing the aircraft to also gather all the necessary inputs for the app every 15 minutes. Therefore, Bluetooth capability was added so that one controller, such as the traffic management coordinator, could collect and input the necessary data to calculate the desired rate. This would then be transmitted to a separate tablet which would display the suggested rate to the second controller. The second controller is notified by a popup that the rate has been updated. In the field trials of 2011, a member of the research team was responsible for gathering and inputting data into the transmitting tablet. The controllers had the receiving tablet, where they had the choice between rate control display and volume control display. In practice, it would most likely be the traffic management coordinator who collects and inputs the data. If gate control is open, this is where the receiver tablet would be placed. If gate control is not open, this position would likely either be covered by clearance delivery or the traffic management coordinator. Figure 7-5 provides a diagram of the Boston Logan control tower. Here you can see the traffic manager coordinator located at the center. From there the information is transmitted via Bluetooth to the gate control position.

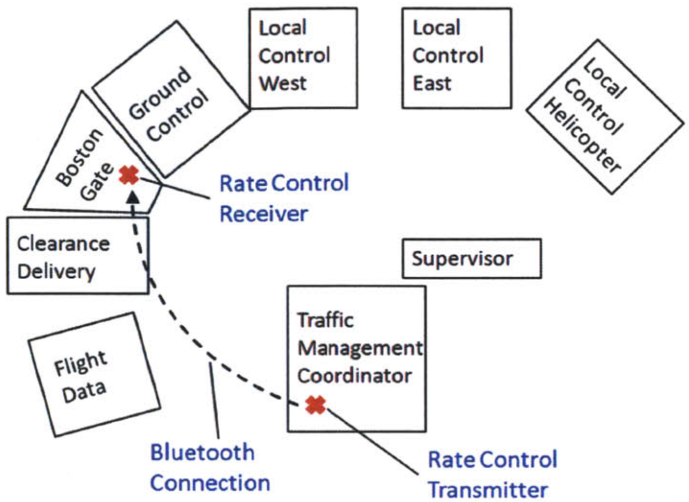

Figure 7-5: Setup in the BOS ATCT

7.4 Analysis of the Implementation

Following the field trials, a survey was conducted of the controllers to get their opinions of the study as a whole, and specifically on the implementation and usage of the tablet. The full survey instrument is included in the Appendix. In this survey, the controllers gave quantitative ratings on five topics: Whether they thought fuel burn decreased, whether surface traffic flows improved, whether throughput was adversely impacted, whether the new (tablet) display was easier to use than the cards used in the 2010 trials, and whether they found the volume control display easy to use. A histogram of the results is shown in Figure 7-6.

It is seen that the survey responses were generally positive, and that the controllers liked the new tablet displays as well. For the questions that were on a positive scale, very few categories were given a rating of 1 or 2, which are the worst ratings. However, there seemed to be some confusion about the one question that was on a reverse scale, where a low rating was actually best. It stated, "Throughput is adversely impacted." It can be hypothesized that there may have been some confusion about the scale on this question of throughput, since the rating of 4 (agree) was correlated with a rating of 4 (agree) on the question of surface flow improvement. Aside from this question, the feedback overall on both the tablet implementation and the congestion control strategy were positive.

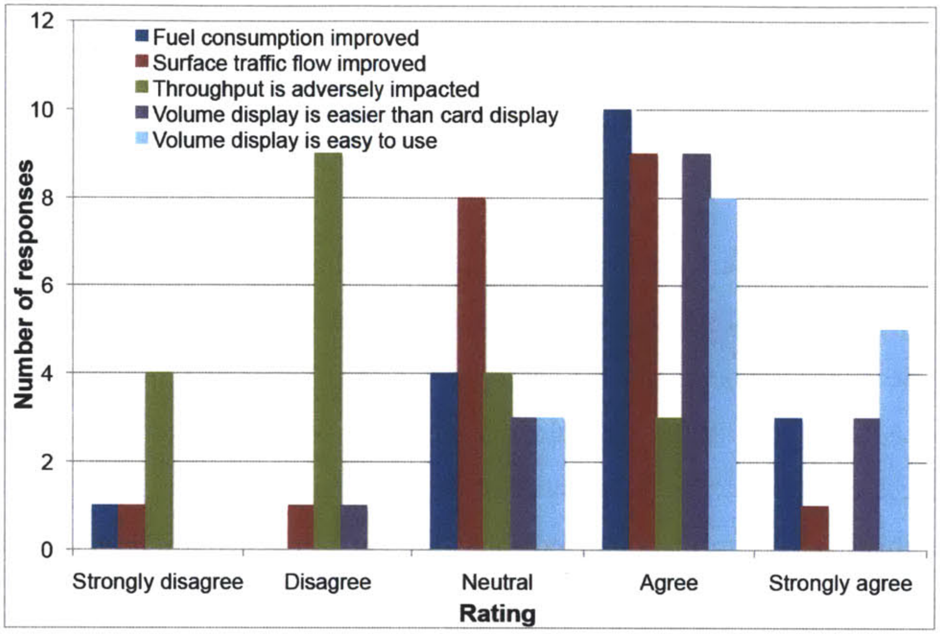

Figure 7-6: Histogram of responses from air traffic controller survey regarding Pushback Rate Control at BOS.

There were some additional open-ended questions on the survey asking the controllers to write in their opinions. For the question of combining two air traffic control positions in the Boston tower, thirteen responses were positive about combining BOS Gate and another position. Ten of these responses suggested BOS Gate should be combined with Clearance Delivery, three indicated it should be combined with the Traffic Management Coordinator, and one person each voted for Ground Control and Flight Data. On the survey it was possible to select more than one possible position for the combination.

For the write-in questions, the results also revealed that the controllers liked the tablet volume control display format. Some of the comments on the best features were highlighting the ability to touch the planes, the ability to reserve spots, and the ability to count the planes and account for aircraft with long delays. One controller liked that the app "allows me to push & tells me to hold". Another said that the app was "easy to use & understand".

For the most part the controllers had no suggestions for improvement on the tablet implementation. The idea that was received for the tablet itself was to increase the aircraft icon sizes on the volume control display. A suggestion for the entire pushback rate control study was maintaining more pressure on the b. In general, the controllers were happy with the modifications made between 2010 and 2011 with one of them remarking, "Liked the improvement in just one year".

7.5 Summary

As can be seen by the comments from air traffic controllers, the tablet implementation of the pushback rate control strategy was well accepted. It is the first step in the strategy of providing decision support to aid the controllers in their duties and optimize the surface operations at the airport. In this case, the researchers or controllers had to manually input certain pieces of data such as the number of departing aircraft under ground control. In the future, this type of information could come directly from a live ASDE-X data feed and eliminate the need for a manual update every 15 minutes. The success of this piece of the strategy encourages further work in the area.

Chapter 8

Conclusions

This thesis has developed a framework for the use of ASDE-X data both for analyzing historical data and for future decision-support tools using ASDE-X. First, the structure and content of the ASDE-X data was described. A process was detailed for extracting only the relevant data for one airport, while also archiving the relevant fields of all the included airports to make future airport analyses more efficient. Then, the process of sorting the aircraft tracks and filling in all the unpopulated fields was described and the remaining steps in the pre-processing phase were listed. An example of an airport visualization was presented with the uses for it described, most notable being the ability to see where queues form. The final processing steps were also listed and some sample outputs were displayed to show some metrics that can be derived from the ASDE-X data.

An analysis was performed on LaGuardia Airport using three months (June-August 2011) of this fully processed ASDE-X data. Statistics were presented at an aggregate level on things like arrival and departure counts and departure taxi times. The next chapter presented a more detailed analysis broken down into the four main runway configurations. For each configuration, the queuing behavior was described, the taxi times were presented, and an analysis of the departure rates and throughput were discussed.

To show how airports differ from one another, a second case study was presented on Philadelphia International Airport. Here the same metrics were derived and discussed on an aggregate level. The two main runway configurations were presented in detail with descriptions of the same analyses. The two airports were compared with each other, first on an aggregate level, and then a comparison was made between the most-used configurations at each airport.

Finally, some next-steps for using ASDE-X data were discussed. This chapter discussed a field trial in which a tablet app was used to extend a previous pushback rate control study in the Boston Logan air traffic control tower. This presents a first step in how real-time ASDE-X data could be used to provide additional decision support to air traffic controllers.

Appendix A

BOS ATC Survey Instrument

Pushback Metering Procedure and Tablet Survey

You are invited to participate in this survey conducted by Ioannis Simaiakis, Lanie Sandberg, John Hansman, Hamsa Balakrishnan and Tom Reynolds from the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology (MIT). The purpose of this survey is to gather feedback on the MIT BOS Metering Demo, and the tablet tool in particular. You were selected as a possible participant in this study because you are an Air Traffic Controller at the Boston facility. Participation in this study is completely voluntary. Survey respondents are encouraged to answer as many of these questions as possible. Surveys are completely anonymous and cannot/ will not be traced back to individual respondents. The results of this study will be possibly included in Ioannis Simaiakis' PhD thesis, or Lanie Sandberg's master's thesis.

Number of years of experience as an air traffic controller:

0-5 years   5-10 years   10-15 years   15-20 years   >20 years

How many years have you worked at the BOS facility?

0-5 years   5-10 years   10-15 years   15-20 years   >20 years

Did you work on the Boston Gate position during the metering test in **2010**?

Yes / No

Did you work on the Boston Gate position during the metering test in **2011**?

Yes / No

Which of the following tools did you use when working Boston Gate during the tests?

Tool A
(physical cards)

Tool B
(rate displayed on tablet)

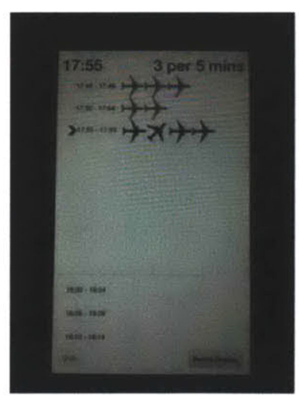

Tool C
(pushback blocks on tablet)

Figure A-1: Controller Survey Page 1

Did you hold any aircraft at the gate during the metering test in **2010**?

Yes / No

Did you hold any aircraft at the gate during the metering test in **2011**?

Yes / No

If you held any aircraft during the test, which of the following tools did you use for doing so?

Tool A
(physical cards)

Tool B
(rate displayed on tablet)

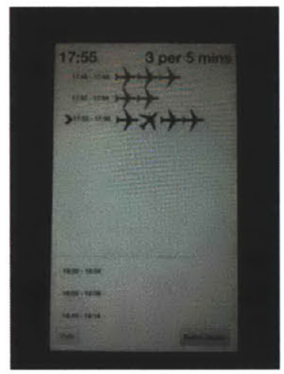

Tool C
(pushback blocks on tablet)

Did you work the Boston Gate position combined with another position during the test of 2011?

Yes / No

If yes, which position(s)?

If the Boston Gate position could be merged with another one, which position(s) do you think would be most appropriate to combine it with?

Figure A-2: Controller Survey Page 2

The metering process saved fuel.				
Strongly disagree	Disagree	Neutral	Agree	Strongly agree
1	2	3	4	5
The metering process improved the flow of aircraft on the surface.				
Strongly disagree	Disagree	Neutral	Agree	Strongly agree
1	2	3	4	5
The metering process adversely impacted the runway throughout.				
Strongly disagree	Disagree	Neutral	Agree	Strongly agree
1	2	3	4	5
The volume interface with pushback blocks (Tool C) made the metering easier to implement compared to the card display only.				
Strongly disagree	Disagree	Neutral	Agree	Strongly agree
1	2	3	4	5
The volume interface with pushback blocks (Tool C) was easy to use.				
Strongly disagree	Disagree	Neutral	Agree	Strongly agree
1	2	3	4	5
What improvements could be made to the process?				
Which features of the tablet did you find most useful?				

Figure A-3: Controller Survey Page 3

Which features of the tablet did you find least useful?
What improvements could be made to the tablet?

Figure A-4: Controller Survey Page 4

Bibliography

[1] Federal Aviation Administration. Aviation system performance metrics database. http://aspm.faa.gov/main/aspm.asp.

[2] Federal Aviation Administration. Fact Sheet Airport Surface Detection Equipment, Model X (ASDE-X). http://www.faa.gov/news/fact_sheets/news_story.cfm?newsId=6296, October 2010.

[3] C. Brinton, C. Provan, S. Lent, T. Prevost, and S. Passmore. Collaborative Departure Queue Management: An Example of Collaborative Decision Making in the United States. 9th USA/Europe Air Traffic Management Research and Development Seminar (ATM2011), June 2011.

[4] Volpe National Transportation Systems Center. Interface Control Document Distribution Of ASDE-X XML Data. Version 1.4, June 2011.

[5] Y. Jung, T. Hoang, J. Montoya, G. Gupta, W. Malik, and L. Tobias. Performance Evaluation of a Surface Traffic Management Tool for Dallas/Fort Worth International Airport. 9th USA/Europe Air Traffic Management Research and Development Seminar (ATM2011), June 2011.

[6] H. Khadilkar. Analysis and Modeling of Airport Surface Operations. Master's thesis, Massachusetts Institute of Technology, 2011.

[7] A. Nakahara, T. Reynolds, T. White, and R. Dunsky. Analysis of a Surface Congestion Management Technique at New York JFK Airport. AIAA Aviation Technology, Integration and Operations (ATIO) Conference, September 2011.

[8] http://www.phl.org/AboutPHL/Pages/Overview.aspx.

[9] T. Reynolds and M. Sandberg. PHL Site Visit Report. Massachusetts Institute of Technology, April 2012.

[10] I. Simaiakis and H. Balakrishnan. Departure Throughput Study for Boston Logan International Airport. 9th USA-Europe Air Traffic Management Research and Development Seminar, January 2011.

[11] I. Simaiakis, H. Balakrishnan, H. Khadilkar, T. Reynolds, R. Hansman, B. Reilly, and S. Urlass. Demonstration of Reduced Airport Congestion Through Pushback Rate Control. 9th USA-Europe Air Traffic Management Research and Development Seminar, 2011.