Analysis of Potential Implementations of Pushback Control at LaGuardia Airport

by

Hector Fornes Martinez

Eng., Technical University of Catalonia (2012)

Submitted to the Engineering Systems Division
and the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degrees of

Master of Science in Technology and Policy

and

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2015

Author: ……
Engineering Systems Division
Department of Aeronautics and Astronautics
January 16, 2015

Certified by: ……
Hamsa Balakrishnan
Associate Professor of Aeronautics and Astronautics
Thesis Supervisor

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

Accepted by: ……...……
Dava J. Newman
Professor of Aeronautics and Astronautics and Engineering Systems
Director, Technology and Policy Program

Analysis of Potential Implementations of Pushback Control at LaGuardia Airport

By

Hector Fornes Martinez

Submitted to the Engineering Systems Division
and the Department of Aeronautics and Astronautics
on January 16, 2015, in partial fulfillment of the
requirements for the degrees of
Master of Science in Technology and Policy
and
Master of Science in Aeronautics and Astronautics

Abstract

Implementations of surface traffic management strategies at congested airports have the potential to yield significant benefits, but must account for the constraints and objectives of multiple stakeholders. This thesis considers the implementation of pushback control policies at LaGuardia Airport in New York. This class of control policies regulate departure pushback rates by holding aircraft at their gates during congested periods, in a manner that maintains the departure throughput of the airport while reducing the taxi-out time. Such a time reduction leads to reductions in fuel burn and emissions. The main contribution of this thesis is the consideration of gate-holding limits at the gate which aim at including operational benefit-cost analysis in addition to the pushback control. The main consequence of those gate holds are gate conflicts and take-off order swaps, which are analyzed in detail throughout this thesis.

The results show that taxi-out savings are a nonlinear and increasing function of gate-holding limit, and thus, more benefits are expected from longer gate conflict limits. However, the non-linear component creates opportunities for additional benefits with marginal cost increases.

On the cost side, gate-holding times are the biggest component, but are commensurate with the benefits. 0ne held minute translates to one saved minute in taxi-out; this finding holds true regardless of the gate-holding limit. Departure order swaps and gate conflicts increase as stricter limits are imposed on the gate-holding times, but not significantly. The benefits and costs are shown to be approximately equivalent to the share of the airlines departures at LaGuardia, demonstrating a fair allocation strategy.

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

ACKNOWLEDGEMENTS

This research was funded by the Federal Aviation Administration.

First, I would like to start thanking my advisor, Prof. Hamsa Balakrishnan, for her support and guidance. Primarily, having the chance of being involved in such an exciting project with Prof. Balakrishnan has made me learn new approaches to frame problems involving complex engineering systems. She has taught me how to use mathematical models to work in multi-stakeholder environment, using quantitative evidence to convince public and private players. I really appreciate her mentorship and the time she has devoted to our meetings and the process of writing academic documents.

Second, I would also like to thank Prof. Richard de Neufville for all his support and mentorship. Prof. de Neufville has taught me very valuable transferable skills that I will definitely apply throughout my professional life. To this end, I am really grateful for his encompassing approach to mentoring.

Third, I would like to thank my wonderful lab mates from ICAT and my colleagues from the Technology and Policy Program, Aero Astro Department and Engineering Systems Division. In particular, I would like to thank Alexandre Jacquillat and Maite Peña, because in addition to being very good friends, they have been my best mentors, providing unconditional professional, academic and personal guidance and support at all times. They have been a source of inspiration during the two and a half years I have spent at MIT.

A warm recognition to my parents Margarita and Francesc, and my brother Carles for their life-long commitment to my education and success. They have been my inspiration during good and tough times during my MIT experience, and with whom I have shared my experiences at MIT.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

Congestion is one of the major challenges currently faced by the U.S. National Airspace System (NAS) and is expected to continue to experience in the foreseeable future. Although only a limited number of airports experience mismatches between supply and demand, (De Neufville et al. 2013) and are therefore congested, these airports have a large impact on the performance of the rest of the network. Delays from these airports propagate to large parts of the system (Pyrgiotis 2012) due to the network and complexity effects. The Joint Economic Committee of the U.S. Senate (Joint Economic Committee 2008) estimated that in 2007, delays cost $40.7 billion to the U.S. economy. The study considered airline operations, passenger cost and other economic activities, and estimated that 20% of the delay costs correspond to the taxi-out phase. These costs include both private and public components; taxi-out delays lead to additional engines-on time, which in turn increases fuel burn for airlines, and emissions (particulate matter, carbon dioxide, hydrocarbons, oxides of nitrogen, oxides of sulfur, among others) – a public cost for the society.

The goal of this thesis is to propose different policies to mitigate the effects of congestion, using New York's LaGuardia Airport as a case study.

1.1. SCOPE

Congestion occurs when the demand for a resource exceeds available capacity. There are broad approaches to managing the problem of congestion at airports: infrastructure expansion, demand management and airport surface management.

First, infrastructure expansions aim at solving the problem by increasing the physical infrastructure (usually additional airport surface and/or additional runways), and hence increasing the airport

capacity. However, there are several concerns with such an approach: first, infrastructure generally requires large amounts of investments; second, there may not be space to expand the airport (usually airports are constrained by the neighboring cities); finally, investing in infrastructure to meet demand may turn into an unsustainable pattern from an environmental perspective.

Demand management, as described in (De Neufville et al. 2013), "refers to set of regulations or other interventions aimed at constraining the demand for access to a busy airfield and/or at modifying the temporal characteristics of such demand". Demand management tries to limit access to an airport through three main measures: overall demand reduction, demand limitations at particular times, and demand shifts from high-demand to low-demand periods. Examples of demand management approaches include schedule coordination (administrative strategy) and congestion pricing (economic strategy), among others. In contrast to infrastructure expansion, demand management aims at reducing demand to capacity levels through access control, instead of increasing capacity to meet demand. Such an approach either denies access or charges a fee for resources (arrival and departure rights).

Third, airport surface management manages the flow of aircraft at the airport during congested periods in order to mitigate the impacts of congestion. It requires a detailed understanding of the queuing processes taking place from the moment aircraft touch down until they take-off. The key difference between surface management and the two previous strategies is that with this approach, there is no change in the physical layout of the airport (capacity is the same) or the flight demand. Thus the performance improvement arises from the way the traffic flows are managed. This approach is the least capital intensive, and the least disruptive of the current everyday airline operations, and therefore, the easiest to implement in practice.

Airport surface management deals with the following processes, which occur after the aircraft has touched-down at the runway:

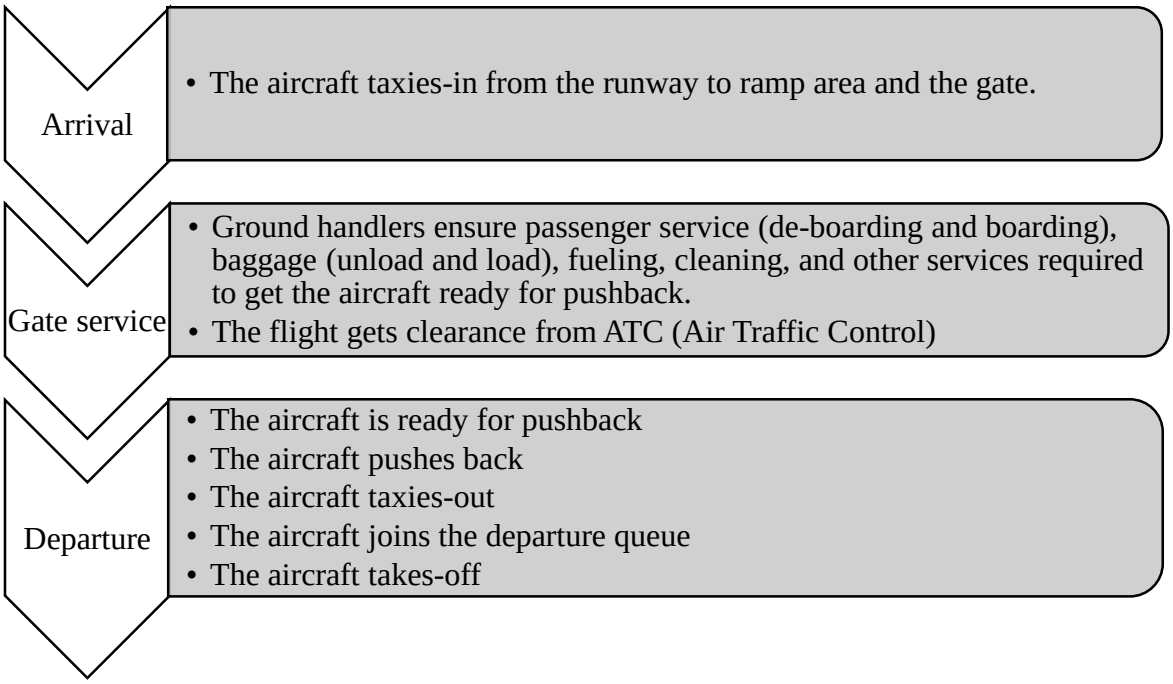

Figure 1: Airport surface processes

This research focuses on the departure part of the process, which corresponds to the last block of processes in Figure 1.

1.2. DEPARTURE METERING-BASED POLICIES: DISCUSSION AND LITERATURE REVIEW

Departure metering is an airport surface management strategy that consists of holding aircraft at the gate to avoid congestion at the runway in periods where the airport experiences saturation. In this regard, Simaiakis (2013) reviews all the most relevant departure management algorithms and classifies them as either trajectory-based models or flow-based (or Eulerian models). While the

former optimizes the individual trajectory of each aircraft, the latter optimizes aircraft counts at different control points; the actuation point in the context of this thesis is the gate, before the pushback procedure. Using flow-based models involves creating virtual queuing approaches, as initially suggested by Feron et al. (1997), and developed by Burgain, Feron, and Clarke (2008). Indeed, the main rationale behind the proposed virtual queues is to ensure the fairness of aircraft queuing without physically queuing at the runway. The virtual queue of concern here occurs at the gate, with aircraft being held before their pushback procedure. A well studied Eulerian approach is the N-control strategy (Pujet, Delcaire, and Feron 2003; Carr et al. 2002; Simaiakis 2009; Simaiakis 2013). In addition to these references, the implementation of the N-control surface management strategy at Boston Logan airport by Sandberg et al. (2014) sets a precedent in the implementation of the proposed strategy.

The main benefits of the N-control metering strategy are the reduction of engines-on time, which in turn leads to a reduction of fuel burn and greenhouse gas emissions. The trade-offs and impact of surface operations to the environment has been analyzed by Simaiakis and Balakrishnan (2009); Simaiakis and Balakrishnan (2010); Ravizza et al. (2013);and Khadilkar (2011), among others.

Building upon the N-control strategy framework developed above, this thesis presents three main contributions:

- The thesis builds a model to implement an N-control surface strategy at LaGuardia Airport, learning from the lessons at Boston Logan Airport. Each airport presents fairly different challenges and constraints that make the implementation of a surface control strategy an ambitious task.
- The thesis proposes maximum gate-holding policies as well as a tool to evaluate these policies during the implementation design process. Metering strategies lead to significant

reductions in taxi-out times; however, pure metering may lead to operational challenges for some stakeholders, especially if the gate-hold durations are high. This thesis develops and evaluates a new portfolio of policies to bring flexibility during the implementation of metering.

- The thesis updates previous studies of metering using RAPT (Runway Availability Prediction Tool) as a weather variable in addition to the visibility conditions (IMC/VMC). RAPT is used as an independent variable in order to predict runway capacity. The statistical analysis confirms the relevance of such an indicator as a runway capacity predicting variable.

1.3. LAGUARDIA AIRPORT

This thesis focuses on analyzing the implementation of an airport surface control management strategy at LaGuardia, one of the most congested airport in the U.S. In order to understand the challenges and opportunities associated with airport surface congestion management, this section presents fixed characteristics of the airport (those that are not likely to change in the short nor in the medium term, such as airport layout, airport terminals, and some stakeholders) as well as "dynamic" characteristics (those that may change in the medium term future, such as airline schedules).

1.3.1. AIRPORT CONGESTION

Figure 2 displays the results of a typical capacity analysis based on the concept of Operational Throughput Envelope (OTE). An OTE is a curve in the 2-D space defined by the average number of arrivals and the average number of departures that can be operated per unit of time. This curve

represents the trade-offs between the arrival throughput and the departure throughput at the airport. Put another way, "the envelopes indicate the capacity that can be achieved for all possible mixes of arrivals and departures" (De Neufville et al. 2013). For the case of LaGuardia, the interrelations of arrivals and departures are particularly relevant given the fact that the two runways intersect and therefore, the capacity of the departing runway is affected by the performance in the arrival runway. OTEs are a function of the runway configuration in use (i.e. set of runways used for arrivals and departures) and the meteorological conditions. In Figure 2, the two OTEs correspond to VMC[1] (Visual Meteorological Conditions) and IMC[2] (Instrument Meteorological Conditions), which, can generally be thought of as "good weather" and "bad weather" conditions, respectively. IMC always generates smaller average throughput than in VMC.

An OTE is generated as follows: First, from a rather large set of historic operational data from departures in a particular period of time (in this case, 15-min periods), mathematical models are used to build envelopes based on runway configuration and visibility conditions (Simaiakis 2013). Second, using scheduling data, for each time period, the number of scheduled arrivals and number of scheduled departures is computed and plotted. This observation may fall outside, inside, or on the edge of the OTE. Figure 2 shows that, at LGA, significant imbalances between demand and capacity may occur, as the scheduling leves fall frequently outside the area defined by the OTEs.

[1] VMC or Visual Meteorological Conditions: Aviation flight category in which visual flight rules (VFR) flight is permitted- that is, conditions in which pilots have sufficient visibility to fly the aircraft maintaining visual separation from terrain and other aircraft. They are the opposite of Instrumental Meteorological Conditions.

[2] IMC or Instrumental Meteorological Conditions: Aviation flight category that describes weather conditions that require pilots to fly primarily by references to instruments, and therefore under Instrument Flight Rules, rather than by outside visual references under Visual Flight Rules.

Note that these imbalances are, of course, more significant in IMC ("bad weather"). These observations justify some intervention aimed at managing surface congestion.

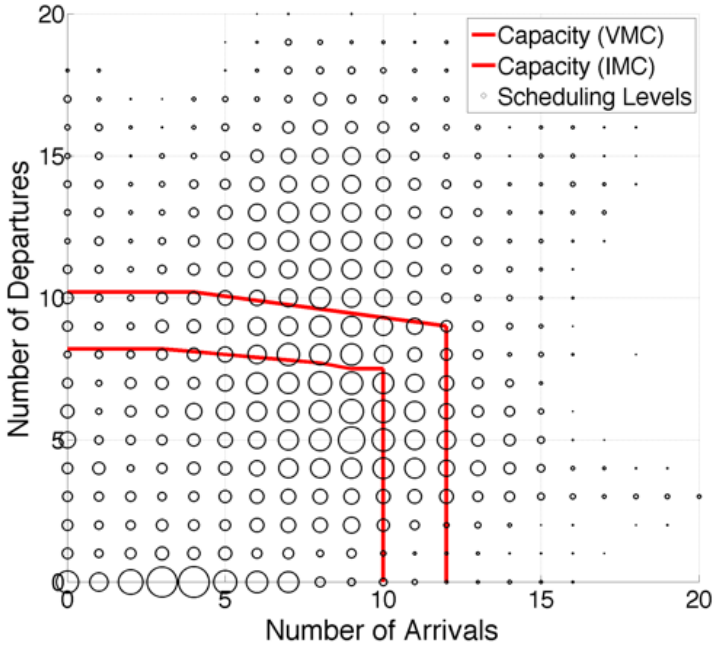

Figure 2: OTEs envelopes at LGA for VMC and IMC, and scheduled number of flights. Source: (Simaiakis 2013; Jacquillat and Odoni 2015)

1.3.2. STAKEHOLDERS

In the context of this thesis there are four main groups of stakeholders:

- Private carriers. This group includes all the airlines operating at the airport; the key among which are Delta, American, US Airways, United, Spirit, Southwest, and JetBlue. Figure 3 depicts the relative number of operations of each airline for the July-August period in 2013. During this Summer period, Delta Airlines, (including all of its regional and shuttle carriers that offer services out of LGA) had nearly a 40% of the departures out of the airport; this number increased during the Fall, Winter, and Spring. Delta's market share is currently

around 50%. The second largest carrier is American Airlines, with approximately a 20% departure share, followed by US Airways, which holds approximately a 10%. The fourth and fifth carriers are United and Southwest with a 5% market share each. In addition to smaller participants such as Spirit or JetBlue, LaGuardia has a 20% of departures operated by other airlines, which include general aviation and charter flights.

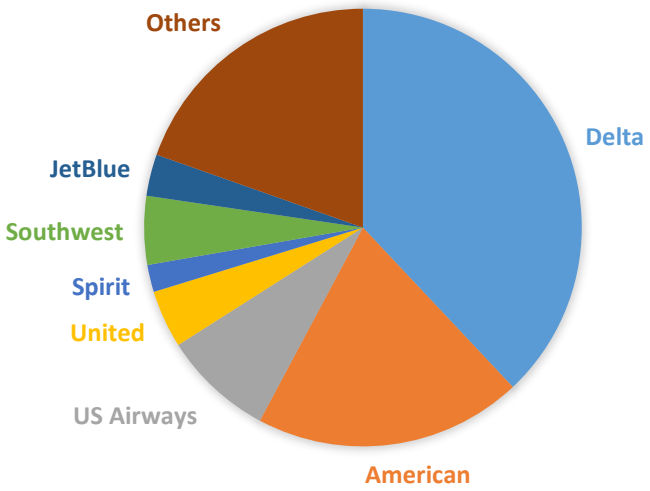

Figure 3: Departure share by carrier at LaGuardia Airport during the July-August 2013 period

- Airport-related institutions. There are two organizations with an essential role at LaGuardia Airport:
    - The Port Authority of New York and New Jersey (PANYNJ) is a Joint organization between the States of New York and New Jersey whose main mandate is to oversee bridges, tunnels, airports and seaports around the New York City area. In the context of this thesis, the PANYNJ is relevant because it oversees LaGuardia Airport, JFK Airport and Newark Airport. Therefore, the PANYNJ has a large stake in any change in airport operations (De Neufville et al. 2013).

- The Federal Aviation Administration is responsible of the Air Traffic Management in the United States, and thus, they decide on which air traffic management policies to implement at all the airports in the country. In this thesis we propose new policies to manage airport surface traffic which needs to be implemented by the FAA controllers at the tower, and thus needs to be approved by the FAA.

- Passengers are important stakeholders in this thesis because they pay the costs of delays, particularly with the lost opportunities.
- Society and the environment are relevant in this problem because the goal is to minimize gas emissions and improve air quality. The former is salient given the significant contribution of air transportation to greenhouse gas emissions, and the latter aims at making to improve the life of people living in the neighborhoods surrounding the airport.

This thesis considers the main trade-offs among the interests of these stakeholders.

1.3.3. LAYOUT

This section introduces the physical layout of the airport, principally the runways, the terminals and the taxiway system. Figure 4 clarifies of the overall layout of the airport. LaGuardia airport has two crossing runways: one runway is oriented in the direction 4/22[3] (S-SW/ N-NE)[4] and the other in the direction 13/31 [5](NW/SE).These two runways intersect and thus, as opposed to what occurs with separated parallel runways, the departing traffic is greatly influenced by the arriving traffic. Put another way, the intersection of runway diminishes the capacity compared with the

[3] Runways are referred to base on their orientation as indicated in (De Neufville et al. 2013)

[4] South-South West/ North-North East

[5] The directions of the runways are displayed in the edges of each runway in Figure 4

situation when these runways can operate independently (De Neufville et al. 2013). However, there are two advantages of having crossing runways: First, they can handle more capacity than a single runway, and thus, in heavily urbanized areas like New York City, they allow more traffic despite the land availability restrictions. Second, they allow arrivals and departures to operate in a variety of weather conditions, particularly different wind directions.

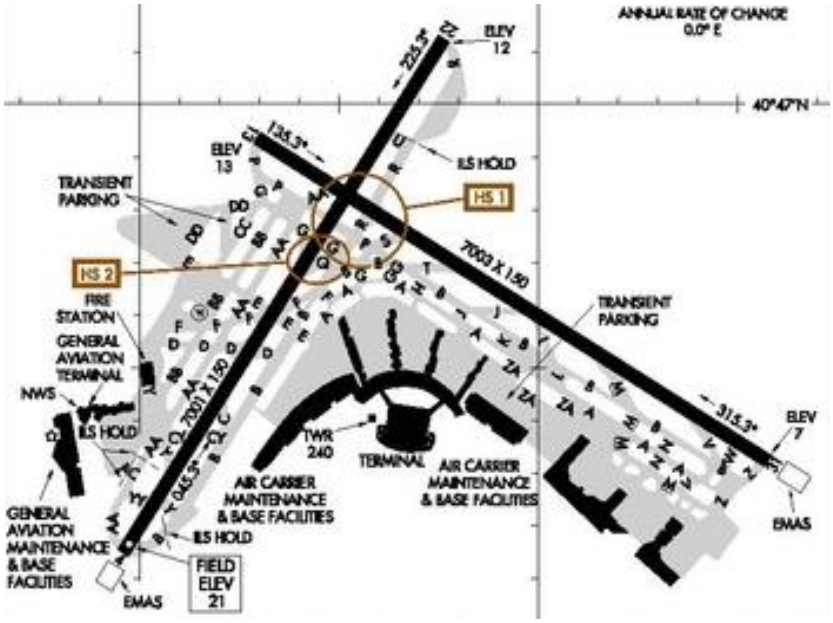

Figure 4: LaGuardia Airport layout, including runways, taxiways and terminals. Source: Federal Aviation Administration, www.faa.gov

The combination of these two runways offers the airport of a portfolio of runway configurations to process the arriving and departing flows. The usual terminology for runway configurations is X/Y, where X is the arrival runway and Y the departing runway. In this regard, the most common runways for the first months of 2013 are displayed in Figure 5. Five configurations (31|4, 22|31, 31|31, 22|13, and 4|13) are the most common in this eight-month period; the rest are not as used much, but having them brings some flexibility for air traffic controllers to adapt airport capacity

to weather issues and noise protected areas in specific periods time. The importance of the five mentioned runways is strengthened in Figure 6, which contains the runway configurations used during the July-August 2013 period, which is the focus time period for this research.

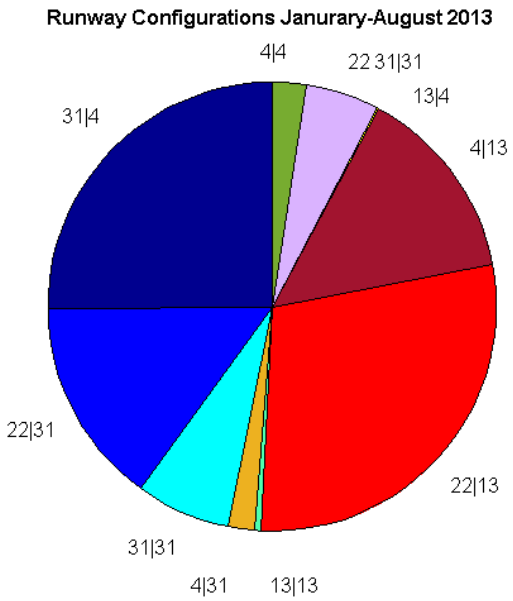

Figure 5: Use of runway configurations during the eight months from January 2013 to August 2013

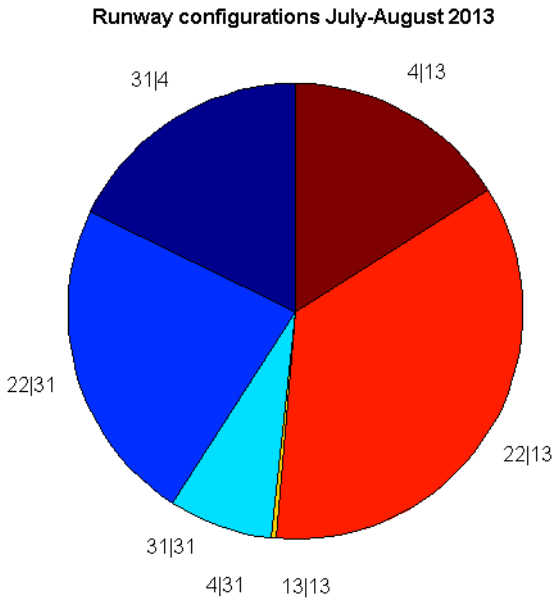

Figure 6: Use of runway configurations during July 2013 and August 2013.

As for terminals, the airport has four terminals: A, B, C and D. In Figure 7 Terminal A is denoted as the Marine Air Terminal, Terminal B as Central Terminal Building, Terminal C as the US Airways Terminal, and Terminal D as the Delta Terminal. Despite the name of the terminals, they do not quite represent the current operations at the airport due to some gate transfers between US Airways and Delta. The main airlines under study operate from the following terminals:

Table 1: List of airlines operating out of each terminals at LaGuardia

Terminal	Airlines
Terminal A	Delta
Terminal B	American, United, Spirit, Southwest, and JetBlue
Terminal C	US Airways and Delta
Terminal D	Delta

Figure 7: Map of the four terminals at LaGuardia Airport. Source: Port Authority of New York and New Jersey, www.panynj.gov

The next airport layout factor to comment on is the taxiway system, which allow aircraft to travel to and from the terminal and the runway. Figure 4 depicts the network of taxiways at LaGuardia Airport.

The airport is constrained by the limited surface available, and that translates to a limited number of taxiways. Such a taxiway system has three important aspects worth noting:

- There is only one taxiway feeding runway 4 | 22 in the southernmost part of the airport, taxiway B. This is particularly salient for departures from runway 4; indeed, in cases of high demand, the departure queue of more than 30 aircraft can use the taxiway up to just before the runway crossing. As for the taxiway AA, parallel to the runway to the left, is mainly used by aircraft operating from Terminal A, in the westernmost area of the airport.
- There are two taxiways parallel to runway 13 | 31, taxiways A and B. These two runways bring flexibility to airport operations, particularly in two circumstances:
    - During periods when the 31 is used as departure runway, these two taxiways are the departure feeders; however, one taxiway is assigned to manage arrivals and the other to departures. By doing so, traffic is not mixed and there is more predictability. Such a strategy should facilitate surface operations without creating additional challenges, and that is why, in case one taxiway is blocked, aircraft may use the other to overcome such blockage.
    - During periods when runways 4, 22, or 13 serve as departure runways, one taxiway is used for eastbound movements and the other for westbound movements. Such a strategy avoids situations with two aircraft facing each other.
- The taxiway density in the eastern area of the intersection is higher than in other areas of the airport, which is helpful at handling departures from runway 13 in periods of high

demand. Taxiing out to runway 13 may be longer and less predictable from aircraft travelling from terminals B, C, and D, because they need to first cross the runway and then wait for their "take-off slot". Such a waiting time requires aircraft to be able to wait in that area without interfering with the arrival traffic. To this end, taxiways, AA, BB, and CC classify aircraft before take-off.

1.4. OUTLINE

This thesis describes the problem of congestion at a constrained airport, focuses on airport surface departure management, proposes a control strategy and, finally, has presentes the characteristics, opportunities and challenges of implementing such a strategy at LaGuardia Airport.

The remainder of this thesis delves deeper into the model and evaluates its performance. Chapter 2 describes the different approaches to airport surface departure control, and derives several control policies based on the strategies described; this description includes mathematical formulation, estimation of parameters, and compilation of data; Chapter 3 presents the performance results of the strategies and policies presented in Chapter 2; finally, Chapter 6 concludes this thesis, summarizing the findings, recommending policies and indicating next research steps.

2. MODEL

This chapter presents the details of the mathematical model built to compare two airport surface control strategies: Pushback-at-discretion (Strategy A), and Metering (Strategy B); the former represents the status quo, and the latter is the proposed strategy that is being evaluated throughout the thesis. Section 2.1 introduces a general overview of the two strategies. Evaluation and assessment of these strategies is done using three datasets: ASPM (Federal Aviation Administration), flightstats (www.flighstats.com), and RAPT (MIT Lincoln Labs). Based on the model requirements, the chapter then introduces the parameters, mathematical tools, and variables that need to be input to the model. One issue that arises when building the model is gate conflicts, which are a consequence of the increase in the time aircraft spend at the gate. Thus, the next section lays out ways to include into the model. Drawing upon all this information, the chapter presents the mathematical formulation associated with each strategy.

2.1. AIRPORT SURFACE CONTROL STRATEGIES FOR DEPARTING TRAFFIC

The strategies below represent two different ways of managing the departing flow aircraft from the moment each aircraft is ready for pushback (in aviation jargon, *off-block time*[6]), until the aircraft takes off (in aviation jargon, *Wheels off-time*[7] or *Take-off time*). As indicated, this section

[6] Off-block time: Time at which the aircraft is ready to start the pushback process. The term *block* refers to the physical objects that are put in front and behind the wheels to prevent de aircraft from moving at all.

7 Wheels-off time: the moment the aircraft is rolling on the runway and the wheels lose physical contact with the runway due to lift. Wheels-off refers to the moment when the wheels are off the runway.

only provides a brief general overview/introduction to these strategies so that the reader can have a general understanding of all the components required for the model.

Before plunging into the description and comparison of the two strategies, it is helpful to introduce a time framework to analyze the processes in a systematic way. Indeed, defining key milestones in the departure process provides a clear pattern for comparison and simulation (as presented later in the thesis). The time framework has four key milestones:

- *EOBT[8]: Earliest Off-Block Time.* Earliest time an aircraft is ready to start the pushback procedure
- *TOBT: Target Off-Block Time.* Time that an aircraft is authorized to start the pushback
- *DQET: Departure Queue Entry Time.* Time when the aircraft joins the physical queue of aircrafts waiting to take-off that starts at the runway heading and growths through the taxiway system. An aircraft joins this queue after taxiing from the gate the queue.
- *ATOT: Actual Take-Off Time.* Wheels-off time for each aircraft

Having this framework in mind, it is possible to introduce and compare the two strategies. In particular, this framework is interesting to look at these other parameters.

- *GHT: Gate-holding Time or Gate Hold.* Time an aircraft is being held at the gate, "prevented" from starting the pushback procedure. Using the framework introduced above, *GHT=TOBT-EOBT.*
- *EOT: Engines On Time.* Total time departing aircraft spends with the engines on burning fuel and emitting gases while on the airport surface. The assumption in this thesis is that

[8] For clarity purpose, EOBT, TOBT, DQET, and ATOT in capital letters will refer to the general variable, whereas when referring to values of these variables for specific flights this thesis uses non capitalized letters: eobt(i), tobt(i), dqet(i), atot(i)

pilots switch on engines slightly before *TOBT* and keep them on after that. Hence, in the context of airport departure surface operations, *EOT= ATOT-TOBT*.

- *TOT: Taxi Out Time.* Time a departing aircraft spends travelling through the taxiway system. This parameter coincides with the time the aircraft has its engines on, and hence *TOT=EOT*.

This framework will allow us to analyze the performance of strategies A and B throughout this thesis.

2.1.1. STRATEGY A: PUSHBACK-AT-DISCRETION

Pushback-at-discretion, from now referred to as Strategy A or the baseline case, represents the way the majority of airports handle the departure processes. The main essence of such a strategy is that aircraft pushback whenever they are ready; put another way: EOBT=TOBT. Given the uncertainty in airport operations (Hall and Fernandes 2013) and the lack of schedule coordination in the U.S. (De Neufville et al. 2013), it is common to see long departing queues at the runway. In other words, with this strategy, when aircraft pushback, they do not know what the current queue length at the runway is, and therefore, after pushing back, they taxi to the runway, join the queue, regardless of the length, and eventually take off. Based on this, it is difficult to set up a straightforward relationship between *DQET and ATOT*. As far as the EOT taxi concerns, it is worth noting that the lack of information on the congestion situation downstream generates useless idling time with the engines on.

In the baseline case aircraft push back at *EOBT* (*EOBT=TOBT*), then, the aircraft taxies out for an uncertain amount of time until they join the queue at time *DQET* and then, after all the aircraft in queue have taken off, the aircraft proceeds to take-off at *ATOT*.

It is clear that information in this strategy propagates only downstream. Indeed, there is no information moving upstream. As a consequence, airlines have an incentive to adopt the following attitude: *the earlier I join the queue, the earlier I will take off.*

2.1.2. STRATEGY B: METERING

Metering consists of holding aircraft at the gate long enough to avoid, as much as possible, the inefficient idling time at the departing queue, without interfering much with the take-off time. In order to do that, this strategy propagates the information upstream, instead of downstream. The model carries out the following process: first, it evaluates the runway departure capacity; second, it propagates this departure capacity through the taxiway system into the terminal, and finally, recommends a pushback rate. Put another way, the model converts runway capacity into "pushback capacity", and aims at limiting the number of aircraft on the surface causing congestion, allowing the pushback procedure only to the number of aircraft that will sustain departure capacity without creating unnecessary congestion. This upstream information process requires two tools to predict departure runway capacity, and the propagation of this runway capacity through the taxiway system. These tools are regression trees and saturation plots and will be addressed in section 2.3.

At this point it is possible to broadly put this proposed strategy into the time framework introduced above. The first difference between the two strategies is that EOBT is not necessarily equal to TOBT due to the definition of gate hold. Indeed, a gate hold is the situation in which an aircraft is "prevented" from starting its pushback procedure despite being ready to do so. It is important to note that gate holds do not occur all the time, and therefore, there may be instances when EOBT and TOBT are still the same. Indeed, gate holds only occur during periods of congestion; in this

regard, section 2.6 explains the details about the conditions in which gate holds occur and the way the model determines TOBT.

As for DQET and ATOT, the difference from Strategy A is that in circumstances when aircrafts are held, the difference between ATOT and DQET is smaller, and so is the EOT, which is the main reason why reductions in emissions and fuel burn exist.

2.2. DATASETS

This research requires two main types of data, weather data and airport scheduling data. As for the meteorological data, the model uses a weather predictability tool called RAPT (DeLaura et al. 2008), which stands for Route Availability Prediction Tool, and was developed by Lincoln Laboratory with the main goal of helping air traffic controllers at airports severely affected by convective weather[9], as is the case of LaGuardia Airport. This tool is effective at analyzing how available a particular air route is on a scale from 0 to 3, where 0 is good weather and 3 corresponds to a route being totally blocked by convective weather and thus, inoperative. In the tower, the numbers are converted into colors and the conversion is the following:

[9] Convective weather: In meteorology, the term is used specifically to describe vertical transport of heat and moisture in the atmosphere, especially by updrafts and downdrafts in an unstable atmosphere. The terms "convection" and "thunderstorms" often are used interchangeably, although thunderstorms are only one form of convection. Cbs, towering cumulus clouds, and ACCAS clouds all are visible forms of convection. However, convection is not always made visible by clouds. Convection which occurs without cloud formation is called dry convection, while the visible convection processes referred to above are forms of moist convection. Source: www.forecast.weather.gov

Table 2: RAPT conversion table from numbers to colors

Number	Color	Weather implications
0	Green	Clear
1	Dark Green	Low Impact
2	Yellow	Caution
3	Red	Blocked

Taking this color code into account, it is possible to understand the RAPT display that air traffic controllers have in the tower, on which they base their decisions. Figure 8 shows such a display, where it is possible to see the different colors available for different routes; such information helps controllers decide upon which route to guide the flow of aircraft through, both arrivals and departures based on the weather situation.

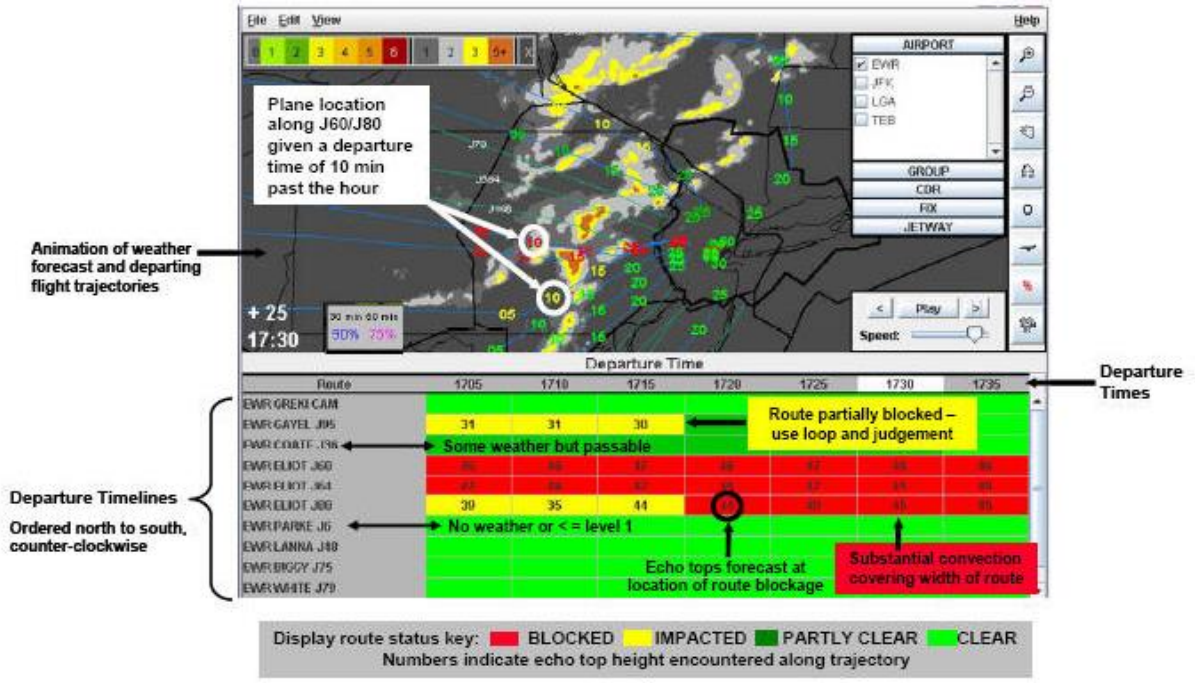

Figure 8: Route Availability Planning Tool display at the tower. Source: (DeLaura et al. 2008)

For the case of this research, the RAPT indicator used is an average all the routes operated to/from the airport. This tool is rather intuitive given its straightforwardness to understand and interpret; however, it is challenging to associate the non-integer numbers (such as 1.67) with the real weather meaning. For the purpose of this research, RAPT is a given variable that is used as a proxy for weather, and as proven throughout the thesis, RAPT is an effective variable for considering meteorological conditions in the context of airport performance.

Regarding the airport scheduling data, this research uses two complementary datasets. On the one hand, the model is fed and most of the parameters are calibrated using the ASPM dataset, which stands for Aviation System Performance Metrics and is developed by the FAA[10]. The most interesting variables from the ASPM dataset are the mentioned below, but the most relevant aspect here is that the information provided is granular, as the data is presented at a flight level, which means that all the parameters are shown flight by flight:

- Airline (Arrivals & Departures)
- Flight number (Arrivals & Departures)
- Runway configuration (Arrivals & Departures)
- Visibility conditions at LGA (Arrivals & Departures)
- Arrival Actual Wheels-on time: Time at which the aircraft put its wheels on the runway.
- Departure Actual call time: Time at which the flight is ready and calls for pushback. In the framework presented above, this time is equivalent to the Earliest Off-Block Time, EOBT.
- Departure Wheels-Off time: Time in the departure roll when the flight lifts off the ground. In the framework presented, this is equivalent to the Actual Take-Off Time, ATOT.

[10] Data search can be done at https://aspm.faa.gov/

It is worth noting that there are two key important data missing: terminal building and gate at which each flight is serviced. This lack of data is what makes the second dataset detrimental for the model; indeed, the model needs flight-specific information on terminal and gate from where the flight is being operated. On the one hand, the terminal information is required to incorporate into the model the pros and cons of operating at each specific terminal, particularly the most salient fact being the difference in taxi-out time that may affect the order dynamics on the surface. On the other hand, the gate information is necessary, when evaluating and avoiding gate conflicts caused by the implementation of the metering strategy. Based on these needs, the second dataset is obtained from the flightstats website[11], which offers flight level details based on generic FAA datasets, but the website buys additional information to airlines in order to provide a better service. In particular, the website adds terminal and gate information, which is of great relevance for this research.

Finally, it is important to clarify the time span of each dataset. Given the way ASPM data is compiled, it is possible to obtain fairly long records of data, and that is the reason why this research uses data from June 2012 until August 2013. However, flightstats data is more cumbersome to obtain as need to be pulled out in 2-hour periods and thus, only strictly necessary data is gathered. This difficulty has some consequences in the parameter estimation as explained in the next section.

2.3. PARAMETERS, MATHEMATICAL TOOLS, AND VARIABLES

This section introduces several parameters, tools and variables that the model requires, and are important to describe before plunging into the model analysis. In particular, this section introduces the concept of unimpeded taxi-out time, decides the length of the time steps used for the simulation,

[11] www.flightstats.com

and finally presents the two prediction tools required to implement metering, characterized by the upstream flow of information.

2.3.1. UNIMPEDED TAXI-OUT TIME

The unimpeded taxi-out time is the taxi-out time that allows the model to simulate a congestion free situation; however, this is a rather broad and unspecific definition. To this end, Simaiakis (2013) carries out a review of different ways unimpeded taxi-out time can be defined. The general definition is *the nominal, free flow taxi out time,* which is related to the absence of obstacles in the taxi-out process. The FAA definition is the following: "taxi-out time under optimal operating conditions, when neither congestion, weather nor other factors delay the aircraft during its movement from the gate to take-off" (Office of Aviation Policy and Plans, Federal Aviation Administration. 2002). From this definition and based on Simaiakis (2009), the following statements can be made:

- The unimpeded taxi-out time is not the minimum taxi-out time; it actually is the average taxi-out time when there is no departure queue.
- Talking about average values is reasonable given that taxi-out times are random variables. Several factors that come into play in this random process are: use of different taxiway routes, use of different speeds, differences in the duration of the pushback process, differences in the process of engine start, or variations in the controller-pilot communications.
- The unimpeded taxi-out time is an average value, which implies than is the result of a calculation, not an observation.

In addition to these observations, it is worth noting that the unimpeded taxi-out time has a weak correlation with the number of aircraft taxiing-out on the surface at the pushback time of each flight (Idris et al. 2001). This is because such an indicator does not consider factors such as aircraft pushing back later but still affecting the taxi-out process of the flight. In order to correct this mismatch, Idris et al. (2001) suggest using the concept of take-off queue for a particular flight, defined as *the number of aircraft taking off between the pushback and take-off time of that flight*; after describing the concept, they also suggest using such a concept to predict taxi-out times.

Based on all these considerations, Simaiakis (2013) sets up the following to estimate the unimpeded taxi-out time. First, he defines *Effective traffic "*for each aircraft l, $N_{eff}(l)$, as the sum if the aircraft taxiing out, $N(l)$, at the time of the flight's pushback t, and the number of aircraft that push back while it is travelling to the departure runway". From a data availability standpoint, $N_{eff}(t)$ requires the model to be fed with the ASDE-X[12] dataset, a more detailed and granular dataset compared to the ASPM the dataset. Clewlow (2010) suggested using another indicator, the *adjusted traffic*, which is equivalent to the effective traffic at being well correlated with the taxi-out time, but it can be obtained from ASPM datasets, which makes the simulations simpler but

[12] ASDE-X enables air traffic controllers to detect potential runway conflicts by providing detailed coverage of movement on runways and taxiways. ASDE-X collects data from a variety of sources to track vehicles and aircraft on the airport movement area and obtain identification information from aircraft transponders. The ASDE-X data comes from surface movement radar located on the air traffic control tower or remote tower, multilateration sensors, ADS-B (Automatic Dependent Surveillance-Broadcast) sensors, the terminal automation system, and aircraft transponders. By fusing the data from these sources, ASDE-X is able to determine the position and identification of aircraft and transponder-equipped vehicles on the airport movement area, as well as aircraft flying within five miles of the airport. Source: www.faa.gov

equally effective. Adjusted traffic is defined f*or "*each aircraft *l*, as the aircraft taxiing out, *N(l),* at the time of its pushback t, and the number of aircraft that push back while aircraft l is taxiing out". The interesting aspect of both the effective and adjusted traffic is the consideration of the traffic ahead while pushing back as well as the traffic that joined while the aircraft is taxiing-out in order to consider all the aircraft that can interfere with this flight's trajectory.

Then the model computes the *empirical unimpeded taxi-out time as the time corresponding to the adjusted (as a proxy for the effective) traffic for which the taxi-out time does not increase, with increasing effective traffic.* To calculate the empiric taxi-out time, it is necessary to use historic data to compute, for each flight, the effective traffic and the observed taxi-out time. Then, for each subcategory of flights, create a scatter plot as seen in Figure 9 with the effective traffic on the x-axis and the taxi-out time in y-axis. For each subcategory we understand all the combinations of factors based on which the results are displayed; in this case results are shown by terminal and runway configuration. For example, the effective traffic and the observed taxi-out time of all flights leaving from terminal B when there is runway configuration 22|31 are depicted in a scatter plot as seen in Figure 9.

These scatter plots allow the model to run a convex optimization regression to fit a non-decreasing, and convex function to the observed data. Such a curve would predict the taxi-out time as a function of the adjusted traffic. As described in (Simaiakis 2013),

> Given *m* pairs of measurements $N_{adj}(l)$ and $\tau(l)$, denoted $(u1,y1), \ldots, (u_m,y_m)$, we seek a convex, non-decreasing function $f_{mean}: \mathbb{R} \rightarrow \mathbb{R}$ that estimates the mean $\tau=f(N_{adj}(l))$. This infinite-dimensional problem is significantly simplified by the fact that N_{adj} is defined only in the domain of natural numbers ($\mathbb{N}_0$). *f* can be restricted to within the domain of $\mathbb{N}_0$ as well, and we need to estimate the values $f(0), f(1), \ldots, f(n)$, where $n=max(N_{adj})$. The function

f is simply a piecewise linear function of *N*, and the monotonicity and convexity constrains are imposed at the points *0,1, ..., max(N_{adj})* by comparing the values and the slopes of a subsequent pieces. *f* is given by the solution to the following convex optimization problem:

$$min \sum_{i=1}^{m} (\hat{y}_i - y_i)^2$$

subject to:

$$\hat{y}_i = f(u_i), i = 1, \dots, m \qquad \text{Eq. 1}$$

$$f(i+1) \geq f(i), i = 0, \dots, (n-1)$$

$$f(i+1) - f(i) \leq f(i) - f(i-1), i = 1, \dots, (n-1)$$

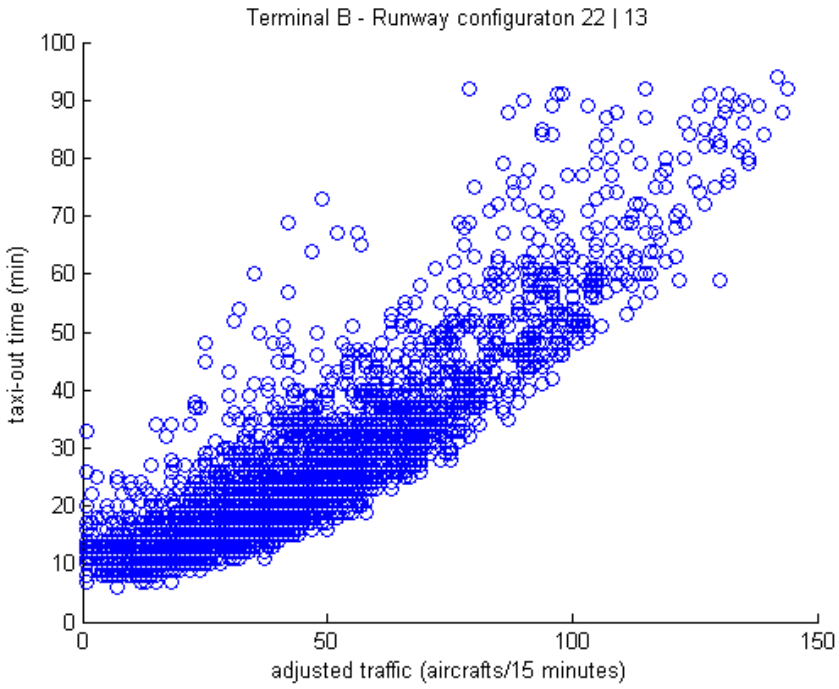

Figure 9: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal B when the runway configuration was 22|13. Data from July 1, 2013 to August 30, 2013

Solving this mathematical problem we obtain the non-decreasing and convex functions, as depicted in Figure 10, as well as the empirical unimpeded taxi-out times *utot(i)* that are a defined

for each runway configuration and terminal. Table 3 displays the empirical unimpeded taxi-out times, which have been obtained by identifying in each non-decreasing, and convex function, the minimum value of the function.

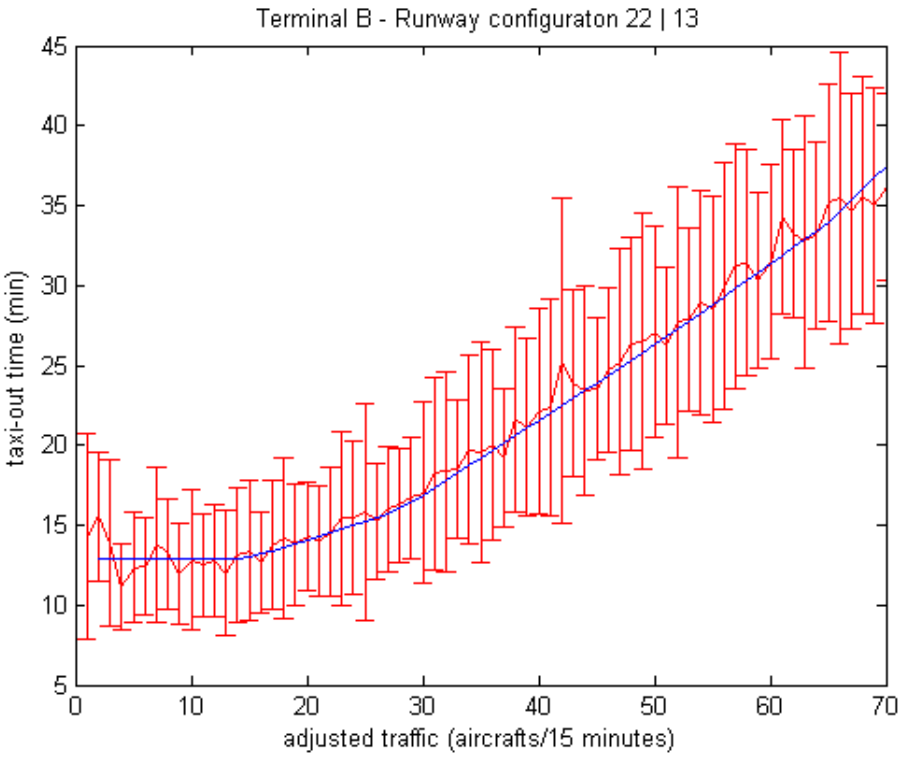

Figure 10: Mean, standard error, and fit function from the scatter plot in Figure 9 for flights from Terminal B when the airport operates with runway configuration 22|13. Data from July 1, 2013 to August 30, 2013

Table 3: Unimpeded taxi-out times (in minutes) for each terminal – runway configuration pair

	31|4	22|31	31|31	4|31	13|13	22|13	4|13	13|4	22 31|31	4|4
T-A	13.92	12.96	12.74	13*	12*	13.6	13*	12*	9.9	12.95
T-B	13.22	10.72	11.39	11*	12*	13.07	14*	12*	14.36	12.41
T-C & T-D	16.63	13.20	12.85	12.66	15*	15.85	15.8	16.5*	10.71	11.89

For clarification purposes, the values with an * have been interpolated from configurations with the same departing runway, given that the pool of observations for that configuration was too small to obtain a reliable enough number.

Another important comment is the reason behind merging Terminal C and Terminal D together. The ASPM dataset does not contain information on terminal or gates; therefore, we need to filter the ASPM dataset by other parameters to indirectly obtain data at a terminal level. Indeed, we identified those airlines and destinations (data available in the ASPM dataset, for each flight) that are served in each terminal, and that provided the model with a de facto terminal based classification. Unfortunately, it was not possible to find a clear classification pattern for Terminals C and D because Delta operates from both terminals and in particular, it serves one same destination from different terminals. This inability to classify flights from terminals C and D is the reason for a combined calculation of unimpeded taxi-out times.

After these two clarifying comments, the results in Table 3 give a sense of the length of typical free-flow taxi-out length; however, it is difficult to compare rows from that table because different terminals have different challenges and different layouts, as can be grasped from the LaGuardia layout in Figure 11. One example of such a challenge is the need for departing flights from Terminal A to cross arrival runway 22 when the departing runway is 31, which is likely to add additional time to the taxi-out time. Another challenge is the availability of handling resources. Recalling that the pushback procedure is one of the key variables affecting the unimpeded taxi-out time, the main carrier in terminal C and D, Delta, has a limited number of resources (tugs and manpower) to serve departing flights. Having more flights to serve with shared resources may affect the efficiency with which Delta carries out pushback. Nevertheless, this is just an example of reasons why comparing rows from Table 3 is challenging. In order to prove the validity of such

an explanation it would be necessary to carry out a resource availability benchmark that would allow us to compare each carrier's performance.

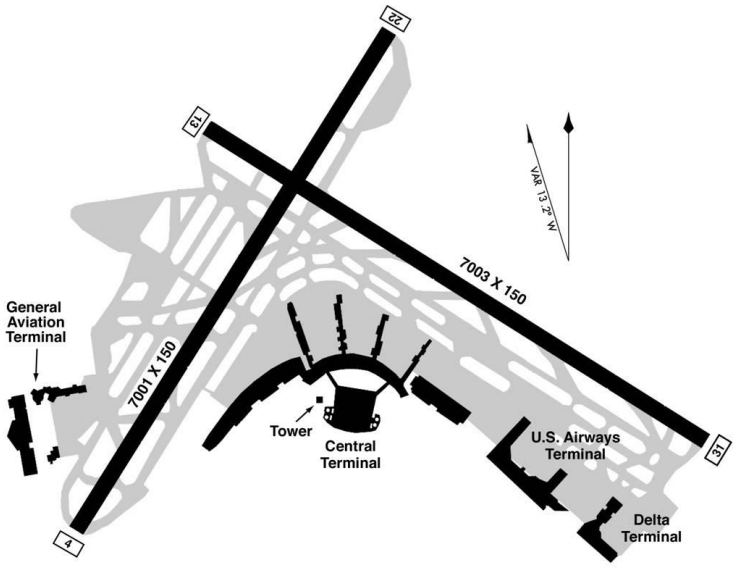

Figure 11: LaGuardia Airport layout, including runways, taxiways and terminals. Source: www.faa.gov

2.3.2. TIME PERIOD LENGTH

At this point we need to decide the length of the simulation time period. This time period is the time fraction in which the model breaks out all days in order to implement a discretized analysis. That is, in order to implement the metering strategy, it is necessary to divide up the whole day into smaller segments. However, choosing the time period length of the simulation is not a straightforward decision. It is, indeed, a key decision and it has different types of implications of which it is worth being aware. The time period length is a trade-off between computational cost, prediction accuracy, usefulness during implementation, and synchronization with the structure of the model and other parameters. These four aspects separately would lead to opposite time period lengths, and therefore, the important exercise here is how to compromise on these opposite trends.

First, from a computational cost point of view, the shorter the time period, the more costly it becomes. Indeed, shorter periods mean more periods and thus, the model needs to carry out a similar routine more time. Therefore, from a computational cost standpoint, the longer the period, the better.

Second, from the prediction accuracy perspective, the shorter the time period, the better. In order to understand such a statement, it is necessary to understand how predictions are carried out. In a nutshell, at the beginning of each period, the algorithm predicts a departure capacity and a pushback rate derived from parameterized tools presented below. From an ongoing work we have seen that the accuracy of these tools becomes less reliable as the length of the time period increases. One particular example of such a trend is the worsening of success rates from the regression trees (presented in section c below), which are being used to predict runway capacity. Table 4 shows that the standard deviation of the error for regression trees increases as the time period length increases, which makes the predictions less reliable.

Table 4: Standard deviation of the prediction error for regression trees with 15-, 30-, and 60-minute time intervals.

Time period length	15-minute	30-minute	60-minute
Standard deviation of the error	1.7 (dep/15-min)	3.6 (dep/30-min)	6.2 (dep/60-min)

Such a worsening trend can be extrapolated to the other tools used at the beginning of each time period. Analyzing the data we can see that the main reason for this behavior is the uncertainty, and particularly the noise. The larger the time period length, the more the noise and thus, the poorer the performance of those tools.

Third, the usefulness during implementation phase has not much to do with the mathematical capability and strengths of the model; it concerns mainly the effectiveness when being implemented in reality. In particular, as indicated in coming sections, the main goal of metering is to suggest a pushback rates to controllers, who then communicate with airlines to give them a prediction of the expected time to start the pushback procedure (TOBTs). One of the strengths of such an approach is the ability to predict, at the beginning of each period, the TOBTs for the duration of the period. This information is greatly appreciated by airlines given that gate-holdings may represent a disruption to them. With such information, airlines can plan their handling resources accordingly. However, the value of this information decays with short notice. That is, information on TOBTs provided 15 minutes in advance is significantly more useful than that information provided 5 minutes in advanced as the ability to plan around it is more salient in the former case than in the latter. Based on such facts, the longer the time period, the more valuable the information provided to airlines is.

Fourth, and last, the synchronization with the structure of the model and other model parameters concerns how this time period lengths fits with the other model components. In particular, given that the model updates predictions and allocations of TOBTs at the beginning of the time period, it is important to match its time length with the dynamics of what happens in that time period. On particularly important factor is the unimpeded taxi-out time, which, as has been shown already, has values around 15 minutes. This makes a very good case to fix the time length to 15 minutes. Indeed, having a time period length similar to the taxi-out time helps isolate the performance of different flights occurring in different time periods. In particular, all the flights that started their pushback during time period T, are likely to join the queue in period T+1; this creates a separation

of flights that is relevant in those cases where there is a change in airport performance conditions in these two periods.

This fourth point justifies the use of a 15-minute time period as opposed to a shorted or longer duration. We believe that 15 minutes is ideal as it leads to a reasonable computational load, it results in rather accurate results, and provides valuable information to airlines to allocate resources based on potential disruptions caused by metering.

2.3.3. SATURATION PLOTS

The next step of the parameter estimation process is to develop one of the two predictive tools that allow the propagation of runway capacity information to pushback information. The tool used for such a task is the saturation plots and its main goal is to measure the capacity of the airport surface. This representation, introduced by Shumsky (1995) and Pujet (1999), presents the surface traffic *N*, on the y-axis, and then the departure throughput *DT*, on the y-axis, as a function of *N*, as depicted in Figure 12.

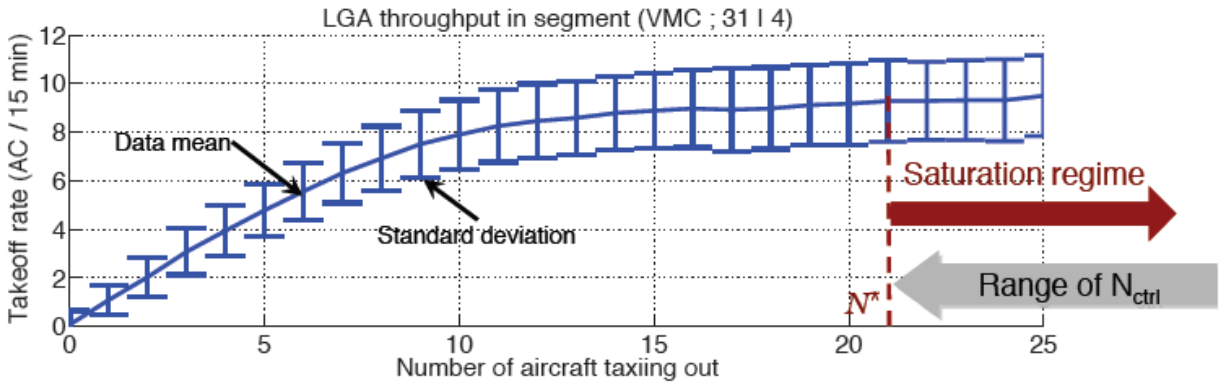

Figure 12: Saturation plot for LGA under VMC rules for runway 31|4. Source: (Simaiakis 2013)

Figure 12 shows that, initially, the departure throughput increases as the surface traffic increases; however, this behavior only occurs up until a saturation point, when demand exceeds the critical value, N*. From Figure 12, the N* value is 21 and the take-off throughput at saturation is approximately 9 aircraft every 15 minutes. This N* can be seen as an intrinsic value of the airport, given that it is operating under a particular runway configuration. Knowing this value, the operator can decide to use it as a control threshold in order to avoid excessive aircraft on the surface. And therein lies the difference between what the literature (Sandberg et al. 2014; Simaiakis et al. 2011) refers to as N-Control and N*. N-control (N_{ctrl}) is the value that the operator may decide to set as a threshold, and N* is an airport configuration-specific characteristic that cannot be modified unless the airport physical characteristics change. This difference leads to the ideal N_{ctrl} discussion. First, N_{ctrl} always has to be larger than N*, otherwise the model would likely be starving the runway, and thus not using the runway capacity efficiently. Second, there is some reluctance to change from the current state which corresponds to N_{ctrl} equal infinity, which is the same as no control of the surface traffic. Such reluctance translates into a pressure to control significantly above the N* value. That is finally a policy decision made by the operator. However, in the context of this research, metering –primarily dependent on the value of N_{ctrl}- is associated with $N_{ctrl} = N^*$. Considering all the above, the model derives the throughput-surface traffic function, fitting a curve to observed data from the ASPM dataset for the first 8 months of year 2013. Indeed, for each 15-minute period, we calculate the surface traffic demand N(t) - defined as the number of aircraft that have pushed back from their gates, but have not yet taken off, and the departure throughput DT(t) - defined as the number of aircraft that take off during the time period of analysis. Then the model implements rather similar to the model used to obtain the unimpeded taxi-out time, except for the concavity instead of the convexity condition. The main goal here is to obtain a monotonically non-

decreasing and concave function that links departure throughput and departure demand (independent variable), based on the observed pairings of *[N(t), DT(t)]*. As presented by Simaiakis (2013):

> The estimation of the data mean regression fit can be formulated as a least-squares problem. Given m pairs of measurements *N(t)* and *DT(t)*, denoted $(u_1,y_1),\ldots,(u_m,y_m)$, we seek a non-decreasing, concave function $f_{mean}: R \rightarrow R$ that estimates the mean $DT= f_{mean}(N)$ *This infinite dimensional problem is significantly simplified by the fact that N is defined only in the domain of natural number (N0).* f_{mean} can be restricted in the domain of N0 as well, and we need to estimate the values $f_{mean}(0), f_{mean}(1), \ldots, f_{mean}(n)$, where $n=max(N_{adj})$. The function f_{mean} is simply a piecewise linear function of *N*, and the monotonicity and concavity constrains are imposed at the points *0,1, ..., max(*N_{adj}*)* by comparing the values and the slopes of a subsequent pieces. f_{mean} is given by the solution to the following convex optimization problem:

$$min \sum_{i=1}^{m} (\hat{y}_i - y_i)^2$$

subject to:

$$\hat{y}_i = f_{mean}(u_i), i = 1, \ldots, m \qquad \text{Eq. 2}$$

$$f_{mean}(i+1) \geq f_{mean}(i), i = 0, \ldots, (n-1)$$

$$f_{mean}(i+1) - f_{mean}(i) \leq f_{mean}(i) - f_{mean}(i-1), i = 1, \ldots, (n-1)$$

A similar process is proposed by Simaiakis (2013) to estimate the median throughput as a function of the departure demand, which can be set up as follows:

$$min \sum_{i=1}^{m} |\hat{y}_i - y_i|$$

subject to:

$$\hat{y}_i = f_{med}(u_i), i = 1, \dots, m \qquad \text{Eq. 3}$$

$$f_{med}(i+1) \geq f_{med}(i), i = 0, \dots, (n-1)$$

$$f_{med}(i+1) - f_{med}(i) \leq f_{med}(i) - f_{med}(i-1), i = 1, \dots, (n-1)$$

An example of a solution for this problem can be seen in Figure 13; the observed data already indicates a saturation trend that the optimized functions follow. In the case of Figure 13, from both the mean- and median-optimized functions we can infer that 23 is the most likely saturation point, or N*.

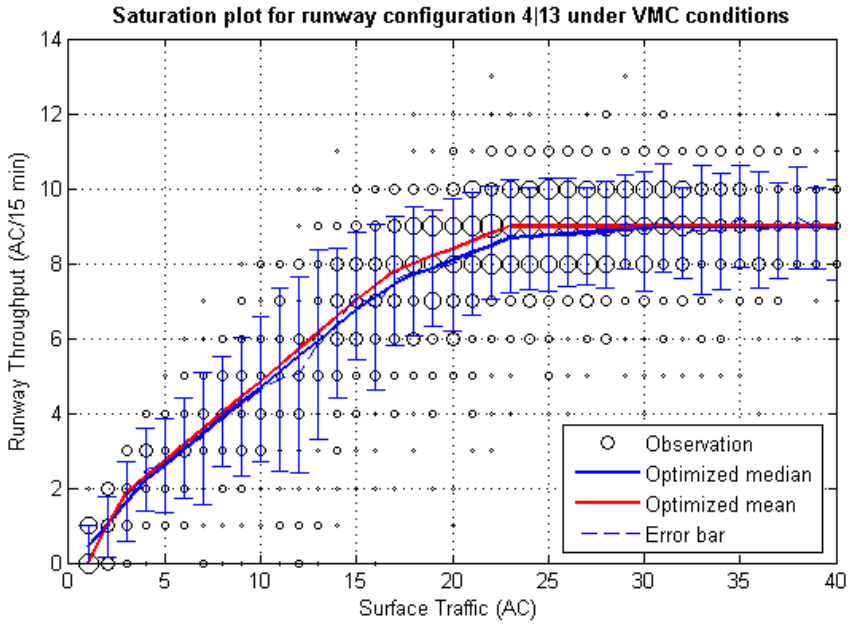

Figure 13: Mean- and median-optimized saturation plots for runway configuration 4|13. The curves are the result of a monotonically non-decreasing and concave function.

It may happen that the mean and the median lead to different N*, which is one of the challenges of the visually recognized N-control strategy. When confronting discerning N* from the mean- and media-optimized functions, leaning towards the smallest or the largest is a matter of how conservative toward the current situation (high N*) or willing to take advantage of the benefits of metering, the decision maker is.

Repeating this procedure for all the runway configurations and for IMC and VMC, it is possible to obtain all the N* values shown in Table 5; these are the values fed into the model to evaluate surface capacity during the simulations.

Table 5: N-control values (# of aircraft) for each terminal – visual conditions pair

	31|4	22|31	31|31	4|31	13|13	22|13	4|13	13|4	22_31|31	4|4
VMC	23	25	18	19	19	20	23	25	25	20
IMC	18	21	16	17	20	21	21	23	25	13

2.3.4. REGRESSION TREES

The second prediction tool that the model requires is the runway capacity regression tree; whose main goal is to predict runway capacity based on observable factors. The predicting variables are arrival rate – which can be obtained from different displays at the tower- and the RAPT value – also available at the tower. It is worth noting that we are building these trees with data from observed saturated periods; that is, the training, validation, and test datasets only include observations where the observed *N* is equal or larger than $N_{ctrl.}$ (which is a function of the runway configuration and the visibility conditions and can be obtained from Table 5). Imposing such a saturation condition ensures that the departure throughput are the highest values, which can be assumed to be the capacity. Indeed, saturation implies that no matter how many more available

aircraft are ready to take-off, there is no more room for these aircraft to do it,; that is the reason why it is assumed that those throughput values correspond to capacity levels.

The process to build regression trees is the following:

- Step 1: Train a classification tree using 70% of the data as training dataset.
- Step 2: Carry out a 10-fold cross validation with 2/3 of the remaining 30% of the dataset to determine the minimum number of leafs (elements of the sample) required on each branch of the tree. For each possible minimum, we calculate the mean ad standard error.
- Step 3: For each minimum number of leafs, carry out a pruning test to determine errors associated with each level.
- Step 4: Build two matrices where the rows are pruning levels and columns are minimum leaf requirements; the first matrix contains mean error values and the second contains standard error values.
- Step 5: Based on these data, decide the minimum requirement and the pruning level that lead to both minimum error and minimum deviation around that error.

The outcome of such a process is a regression tree like the one shown in Figure 14. The same process needs to be carried out for every configuration under VMC and IMC conditions. A quick interpretation of the tree in Figure 14 is: if the Arrival Rate is larger than 5.5, the RAPT larger than 0.5, and the Arrival Rate larger than 6.5, then, the expected runway capacity is 5 aircraft every 15 minutes.

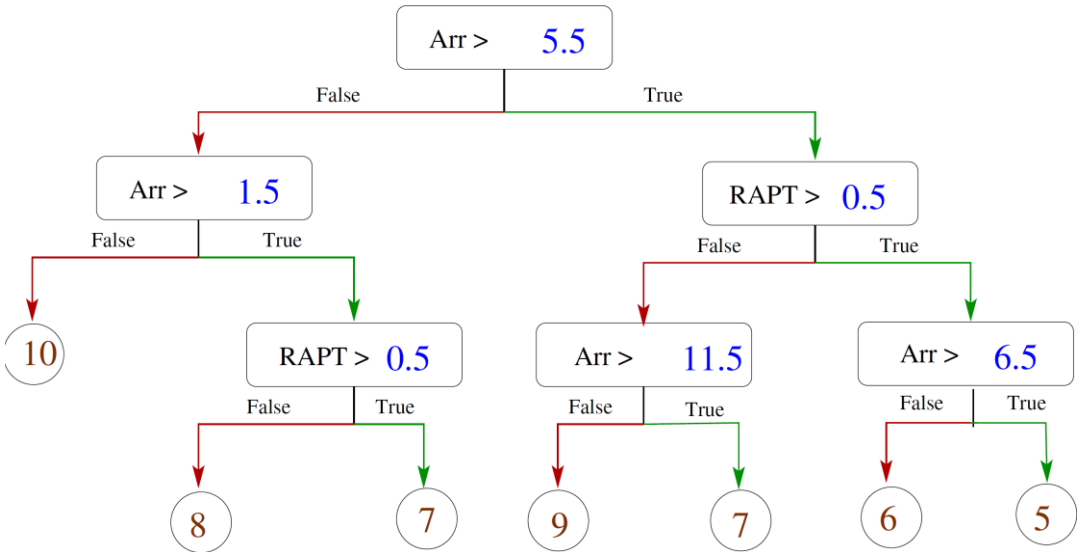

Figure 14: Regression tree for runway configuration 4|4

2.3.5. DEPARTURE SLOTS

For each simulation period T, the model generates a *Take-off Slot Vector, TOSV(T)*, which contains all the *Take-off Slots* for each simulation period T, or $tos_T(k)$, where k=1,2,… n_T. n_T is the number of take-off slots [13]in time period T, and is equal to the number of take-off operations observed in the ASPM dataset in that time period T. The rationale is that given that the model is not trying to alter the nature of the take-off process, the take-off throughput is expected to be the same in the experienced and in the simulated case. After obtaining n_T, the model homogeneously distributes n_T slots over time period T, with an average time separation between tos_T of T/ n_T.

[13] It is worth noting that both vectors represent internal tools of the model, which are used to simulate the departure capacity. When talking about slots, this thesis is not referring to the administrative system to manage departure demand in most of the airports outside the United States. For more information, see (De Neufville et al. 2013)

$$TOSV(T) = \begin{pmatrix} tos_T(1) \\ tos_T(2) \\ \vdots \\ tos_T(n_T) \end{pmatrix} \qquad \text{Eq. 4}$$

After calculating the take-off slots for all the time periods, it is necessary[14] to put all of them together in a "day-long take-off vector" that includes all the TOSV(T) on day *d.* Such a vector is referred to as *Day Take-Off Slot Vector, or DTOSV* and is built as follows:

$$DTOSV = \begin{pmatrix} TOSV(1) \\ TOSV(2) \\ \vdots \\ TOSV(N_T) \end{pmatrix} = \begin{pmatrix} \begin{pmatrix} tos_1(1) \\ tos_1(2) \\ \vdots \\ tos_1(n_1) \end{pmatrix} \\ \begin{pmatrix} tos_2(1) \\ tos_2(2) \\ \vdots \\ tos_2(n_2) \end{pmatrix} \\ \vdots \\ \begin{pmatrix} tos_{N_T}(1) \\ tos_{N_T}(2) \\ \vdots \\ tos_{N_T}(n_{N_T}) \end{pmatrix} \end{pmatrix} = \begin{pmatrix} tos(1) \\ tos(2) \\ \vdots \\ tos(N_{add}) \end{pmatrix} \qquad \text{Eq. 5}$$

For the sake of clarification, N_T is the number of time periods in one day and N_{add} is the length of all the take-off slots.

Associated with the DTOSV, the model creates a *Take-Off Slot Availability Vector, or TOSAV.* This is a binary vector (0 or 1) that indicates, for each take-off slot tos(k), whether the slot is available (takes value of 0) or used (takes value of 1).

[14] It is necessary because aircraft do not necessarily pushback, taxi and take-off in the same time period due to the length of the period and congestion. Put another way, an aircraft that is ready to take-off in one time period may have aircraft ahead that prevent that aircraft from taking-off until some time periods later.

$$TOSAV = \begin{pmatrix} tosa(1) \\ tosa(2) \\ \vdots \\ tosa(N_{add}) \end{pmatrix} \quad \text{Eq. 6}$$

While doing the simulations, the algorithm needs to update each *take-off slot availability*, tosa(k), from 0 to 1, when the take-off slot tos(k) is allocated to a particular flight.

2.4. GATE CONFLICTS

The metering strategy has several advantages that make its implementation interesting. However, the nature of metering, which consists of holding aircraft at the gate for a longer time, is likely to cause additional gate conflicts. A gate conflict is defined as the situation in which an arriving aircrafts touches-down (wheels of the landing gear touch the runway) at a time when the gate the aircraft is assigned to is still being used by the previous aircraft. These conflicts are the consequence of the late departure of the departing aircraft and/or an early arrival of the arriving aircrafts. As for metering, holding aircraft at the gate increases the chances of "next flights" touching-down while the departing flights are being held.

Several reasons can increase the likelihood of gate conflicts, among which it is worth noting two: firstly, those airports or terminal areas where gates are a scarce resource are more likely to experience gate conflicts as there is not much room to find alternative gates; secondly, airlines that adopt an aggressive approach to scheduling (they operate a lot of flights out of the same gate) are more likely to experience gate conflicts. The latter reason may look interesting as a way to ensure earlier releases in a metering environment; this would be true if airports behaved in a rather deterministic manner. However, airport operations are exposed to a high degree of stochasticity that is likely to play against the interest of the airlines. Indeed, aggressive scheduling is likely to lead to long waiting time for arriving aircrafts that have no alternative gate to use and need to wait,

blocking the way for other aircraft. Put another way, airlines have no objective incentive to game the system by provoking gate conflicts, because such a strategy is expected to negatively affect the performance of airline operations at the airport. However, Kim and Feron (2014) proved that airlines can reduce the effect of gate conflicts caused by metering, through the implementation of a more robust gate assignment; they recommend a robust gate assignment together with metering to improve the performance at congested airports.

2.5. STRATEGY A: PUSHBACK-AT-DISCRETION

As indicated by the name of the strategy and as previously described, pushback-at-discretion implies no control over the pushback procedure, which translates into a downstream information process; this section introduces the mathematical formulation of this strategy. In order to keep consistency and clarity throughout the thesis, this section strictly follows the time framework presented in section 2.1, applied to each particular flight. That is, for each flight, the algorithm needs to be able to simulate eobt(i), tobt(i), dqet(i), and atot(i).

2.5.1. MATHEMATICAL FORMULATION

Taking all the above into consideration, the following presents all the steps to model the pushback-at-discretion strategy for one particular day. Carrying out all the steps subsequently allow the model to simulate as many days as required.

Step 1: Create a matrix with as many rows as flights in the simulated day. Column 1 will be filled out with EOBTs, column 2 with TOBTs, column 3 with DQETs and column 4 with TOTs. Column 5 contains information on the terminal where each flight starts its taxi-out process, column 6 presents information on visibility conditions (IMC, VMC), and column 7 has information about

runway the runway configuration. The data in both column 6 and 7 correspond to the conditions (visual conditions and runway configuration) at the time when each flight starts its push-back procedure.

Step 2: Fill out EOBTs and TOBTs. The algorithm assumes eobt(i) and tobt(i)to be the "actual pushback time" from the ASPM dataset, which is the time that flight i started the pushback process. The rationale behind this assumption is that in the current strategy aircraft tend to start the pushback process as soon as they are ready. Thus, the eobt(i) – time when flight i is ready to push back- and the tobt(i) – time when flight i actually pushes back – coincide and are equal to the time the plane pushed back in reality, which corresponds to the "actual pushback time".

Step 3: Evaluate the unimpeded taxi-out time for each flight, utot(i), as a function of the terminal, the runway configuration and the visibility conditions.

$$utot(i) = f(tb, r, v) \quad \text{Eq. 7}$$

tb: terminal building

r: runway configuration

v: visibility conditions

Step 4: Fill out DQETs. the mathematical expression for dqet(i) can be obtained as follows:

$$dqet(i) = tobt(i) + utot(i) \quad \text{Eq. 8}$$

The assumption here is that aircraft circulate through the taxiway system for a time equal to the unimpeded taxi-out time (Travel time in Figure 15) until they join the departure queue. This assumption is simple but is effective at including important factors such as the difference in taxi-out time from different terminals that leads to changes in the queue order. Figure 15 divides the taxi-out time in the travel portion (Module 1) and the queuing portion (Module 2). Module 1 in

Figure 15 effectively shows that travel portion of the taxi-out process is only affected by the ramp delays and taxiway delays, which are mainly dependent on the terminal building of origin.

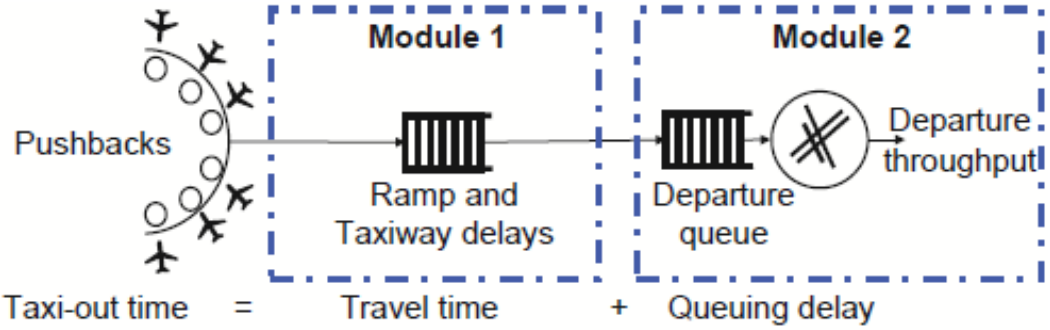

Figure 15: Taxi-out framework that divides the taxi-out time in travel time and queuing delay. Source: (Simaiakis 2013)

It is also worth noting that, DQET is related to the real performance at the airport; it is actually the end of the travel portion following the framework from Figure 15 , right before the aircraft starts idling. The main deviation from reality is the uncertain nature of the unimpeded taxi-out time, which makes it the most likely value in absence of surface traffic. Put another way, the unimpeded taxi-out time is a useful mathematical tool to tool to calculate ATOT, although it might be difficult to clearly observe it when compared with other milestones.

Step 5: Sort all flights by DQET, which provides a list of the flights based on the order they are ready off take-off, after having taxied through the taxiway system and joined the departure queue at DQET.

Step 6: Assign a take-off slot atot(i) to all the flights following the order of joining the queue. For each flight i proceed to do the following sub-steps:

Step 6.1: Identify the earliest available take-off slot after dqet(i). That is, the slot should minimize the time difference between the slot and the dqet(i), verifying that the availability

vector TOSAV has a value of 0 for that slot. It is worth noting that here the flight is fixed and the variable the model is seeking to allocate is the take-off slot, tos(k), from the DTOSV vector. Mathematically, this condition could be expressed as looking for k such that:

$$min|\ tos(k) - dqet(i)\ |\quad k = 1,2,\dots,N_{add}$$

subject to:

$$tos(k) \geq dqet(i)\quad k = 1,2,\dots,N_{add} \qquad \text{Eq. 9}$$

$$tosa(k) = 0\quad k = 1,2,\dots,N_{add}$$

Many tools provide solutions to this problem. Given that the simulations in this thesis are run in Matlab, the tool used to solve the take-off slot search is the Matlab *find* function.

Step 6.2: Update the TOSAV. After allocating a take-off slot to a flight, the corresponding component of the availability vector TOSAV should change to 1.

Step 6.3: Repeat sub-steps 6.1 and 6.2. for all the flights in the day.

2.6. STRATEGY B: METERING

Metering requires a more complex mathematical formulation to account for the upstream propagation of information that will allow a reduction in the surface traffic. Using a similar approach as for Strategy A, the goal here is to simulate eobt(i), tobt(i), dqet(i), and atot(i) for each flight i. One new feature in the formulation of Strategy B is the need to carry out the simulation in time steps. By segmenting the time, the model is able to improve performance.

2.6.1. MATHEMATICAL FORMULATION

Before plunging into the step by step formulation of the model, it is worth having a general understanding of how the model goes about propagating the information upstream. It is important to recall the following: the main objective is to be able to predict runway capacity data, and then propagate these data upstream to suggest a pushback rate that limits the surface traffic to those aircraft that strictly ensure number of aircraft that get to the runway and need to idle before they can take-off. A three-stage process provides a clear context of the proposed strategy built in this thesis:

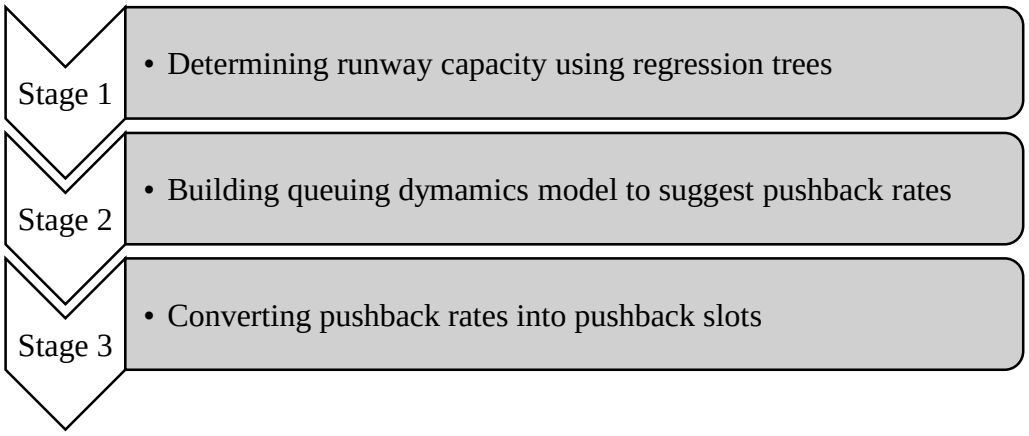

Figure 16: Conceptual stages to describe the metering strategy

- **Stage 1** consists of determining runway capacity using the regression trees already introduced. That is, based on the runway configuration, the visibility conditions, the rate of arrivals and weather (RAPT), the trees provide a predicted value of runway capacity. As indicated previously, it is important to use predicted capacity (from the regression trees) instead of observed throughput in those situations when we are forecasting instead of analyzing "what would have happened if". In this case, the goal is to predict the capacity

before this capacity is experienced, and thus, it would be conceptually inappropriate to use the observed capacity as it has not been observed in the time when the prediction is made.

- **Stage 2** primarily intertwines all the variables involved in the queuing dynamics at the taxiway level. In this stage, the main goal is to understand the dynamics of the ramp area and taxiway system, from section 2.3, in order to evaluate how to translate the runway take-off capacity translates to surface traffic. In order to do so, the model imposes the conservation principle the taxiway system, taking into account inputs, and outputs. That is, in a particular time period, there needs to be control over the number of aircraft that leave the system, those that stay. All these variables need to be aligned with the saturation considerations coming from the saturation plots presented above. By imposing the no-saturation condition it is possible to obtain the number of aircraft that need to push back in an easy expression. Figure 17 provides a framework to deduce such an expression:

$$N_{push} = N_{ctrl} + N_{takeoff} - N_{cur} \qquad \text{Eq. 10}$$

where N_{push} is the number of aircraft that need to push in order to satisfy the no-saturation condition; N_{ctrl} is the saturation threshold from the saturation plots above; $N_{takeoff}$ is the expected take-off capacity obtained in Stage 1; and N_{cur} is the number of aircraft that are on the surface at the beginning of the time period.

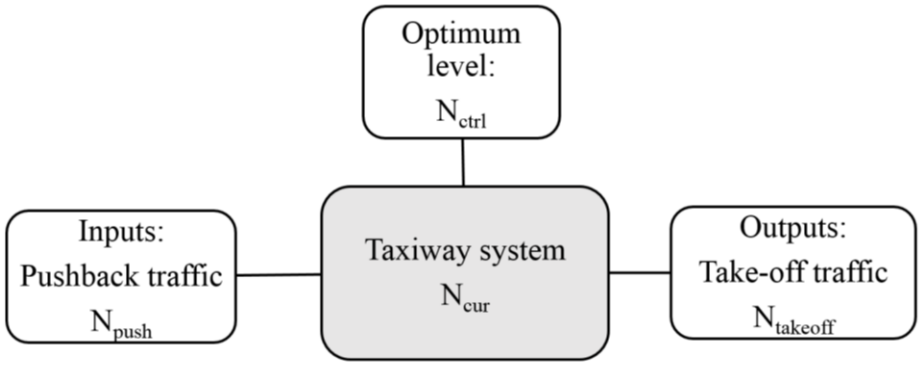

Figure 17: Taxiway system queuing model to deduce N_{push}

Stage 3 translates N_{push} to a pushback slot vector PBSV(T) equivalent to the take-off slot vector TOSV(T) derived in section 2.3. Indeed, the model homogeneously distributes N_{push} pushback slots (pbs) over time period T, with an average time separation between pbs_T of T/ N_{push} .

$$PBSV(T) = \begin{pmatrix} pbs_T(1) \\ pbs_T(2) \\ \vdots \\ pbs_T(n_T) \end{pmatrix} \quad \text{Eq. 11}$$

Similar to what happened with the take-off slots, each PBSV has a pushback slot availability vector that indicates whether a particular slot has been used or not in a particular moment.

$$PBAV(T) = \begin{pmatrix} pbsa_T(1) \\ pbsa_T(2) \\ \vdots \\ pbsa_T(n_T) \end{pmatrix} \quad \text{Eq. 12}$$

With this strategy, it is possible to limit the number of aircraft in the departure queue and thus reduce the queuing time. It is interesting at this point to recall the framework presented in Figure 15. With the metering strategy, the taxi-out portion of queuing delay is being reduced significantly and thus, the taxi-out time gets significantly limited to travel time.

This metering framework has several expected benefits, but it also important to consider the main challenges associated with it; the most important of these challenges being gate conflicts. In this fashion, Figure 18 depicts the simulation process implemented in this research in order to consider these gate conflicts. First, the model carries out a simulation with a pure metering strategy; understanding pure metering as a strategy where there are no constraints and therefore aircraft can be held for as long as they need to be in order to avoid unnecessary queuing delay. Second, the model analysis the gate conflicts by looking at the interactions between arrivals and departures out of each gate, for all the gates at the airport; by doing this it is possible to identify flights for which

the holding time leads to a conflict with an arriving flight. Third, the model identifies these conflicting flights and modifies their pushback times in order to avoid this situation.

Figure 18: New simulation process to include gate conflicts

The remainder of this section describes in detail the steps to implement this three-process simulation framework. It is important to notice that although the format of the steps may seem different from the steps in Strategy A, they are analogous and the their target is the same. Indeed, the two sets of steps aim at deriving EOBTs, TOBTs, DQETs, and ATOTs for all the flights.

Table 6: Steps to simulate the departure procedure including gate conflict analysis and the correction to include those conflicts in the model

Process	**Steps**
Simulation 1	**Step 1**: eobt(i) for all flights is equal to the actual pushback time from the ASPM dataset **Step 2:** Update the time counter. T = T+1 **Step 3:** Compute N_{cur} (Number of aircraft currently active on the surface) **Step 4:** Evaluate N_{ctrl} as a fuction of the runway confuguration and the visibility conditions $N_{ctrl}=f(r,v)$ **Step 5:** Compare N_{ctrl} and N_{cur}

	If $N_{cur} < N_{ctrl}$: Follow the steps in strategy A, only for the flights in the time period under analysis. If N_{cur}>Nctrl : Go to step 6 **Step 6:** For all flights from previous time periods without a pushback slot, the algorithm allocates the pushback slots available on a FCFS basis. If pushback slots are still available they are allocated to the flights calling for pushback in the time period analyzed in a FCFS manner. Allocating a pushback slot to a flight means searching for the first $pbs_T(k)$ available after eobt(i); put another way, the algorithm imposes tobt(i) equal to the $pbs_T(k)$ which fulfills the availability conditions. Those flights without pushback slot allocated have to wait until subsequent time periods. **Step 7:** For all flights that have a tobt(i) assigned, dqet(i) is deduced as in Strategy A: $dqet(i) = tobt(i) + utot(i)$, with utot(i) as a function of the terminal, the runway configuration and the visibility conditions. **Step 8:** For all flights with dqet(i) calculated, atot(i) is obtained by following Step 6 and sub-Steps 6.1, 6.2 and 6.3 from Strategy A. In brief, it is necessary to identify, for each flight, the earliest take-off slot available and then update the availability vector.
Gate conflict analysis	**Step 1:** List all the arriving and departing flights from the ASPM dataset **Step 2:** List all the gates for each arriving and departing flight from the flightstats dataset. **Step 3:** For each gate, create a 3 column matrix with all the departing and arriving aircraft that are serviced at that particular gate. The first column shows with a -1

	or a +1 whether the flight is a departure or an arrival. The second column contains the TOBTs of departing aircraft and the wheels-on time of arriving aircraft. The third column shows the gate-holding length due to metering. **Step 4:** Determine those departing flights for which the next arriving aircraft touches down during the additional gate-holding. Flights that fall in this category are the ones that need to pushback soon after an arriving aircraft touches down. **Step 5:** Create a "gate conflicts matrix" with 2 columns with length equal to the number of flights in the simulated day. The first column contains those departing flights that need to push back earlier as they are involved in gate conflicts. Flights with such a condition are displayed with the number 1; otherwise, this column is 0. For those flights with 1 in the first column, the second column provides the time when the conflicting aircraft needs to start the pushback procedure, which corresponds to the time the next aircraft touches down.
Simulation 2	**Step 1:** From the list provided by the gate conflicts analysis, identify all the flights that are affected. **Step 2:** For those flights, change the tobt(i) to the value shown in the second column of the gate conflicts matrix. **Step 3:** For all flights repeat Steps 4, 5 and 6 from Strategy A to allocate the new dqet(i) and atot(i).

2.7. DEFINITION OF MAXIMUM GATE-HOLDING POLICIES

After describing all the details of the metering strategy, the natural transition would be to proceed to show results; however, at this point it is best to run a first round of simulations and see how the model is performing. In particular, we would like to look at the most controversial output of the model, the gate-holding times. Indeed, the gate-holding times are what ensures the efficiency of the model, but are also perceived as an operational challenge from an airline's perspective. To this end, Figure 19 plots the duration of the gate holds at a flight level, a consequence of the metering strategy. Three main conclusions can be drawn from Figure 19:

- Between 80% and 90% of the flights are not being held at all
- The frequencies associated with gate holds longer than 10 minutes are small
- The tails of the distributions are rather long, but with low frequency.

We anticipate that the last point may bring some discussion among stakeholders. Indeed, from a practical perspective, managing a 5-minute gate hold may be manageable by an airline current operation procedures given the current uncertainty. However, a 20- or 25-minute gate hold may be an operational challenge. Given this reality, this research proposes a portfolio of three metering policies to help in order to provide some flexibility in the implementation of the metering strategy.

- Policy B.Unr: This policy corresponds to a pure metering strategy as described so far with no limit to the gate hold length.
- Policy B.15: This policy is limits the gate-holding time to 15 minutes. That is, after 15 minutes of holding, the aircraft is allowed to start the pushback procedure even though from a metering standpoint it would be recommended to wait.
- Policy B.10: This policy is similar to the previous one, but in this case the gate-holding limit is 10 minutes.

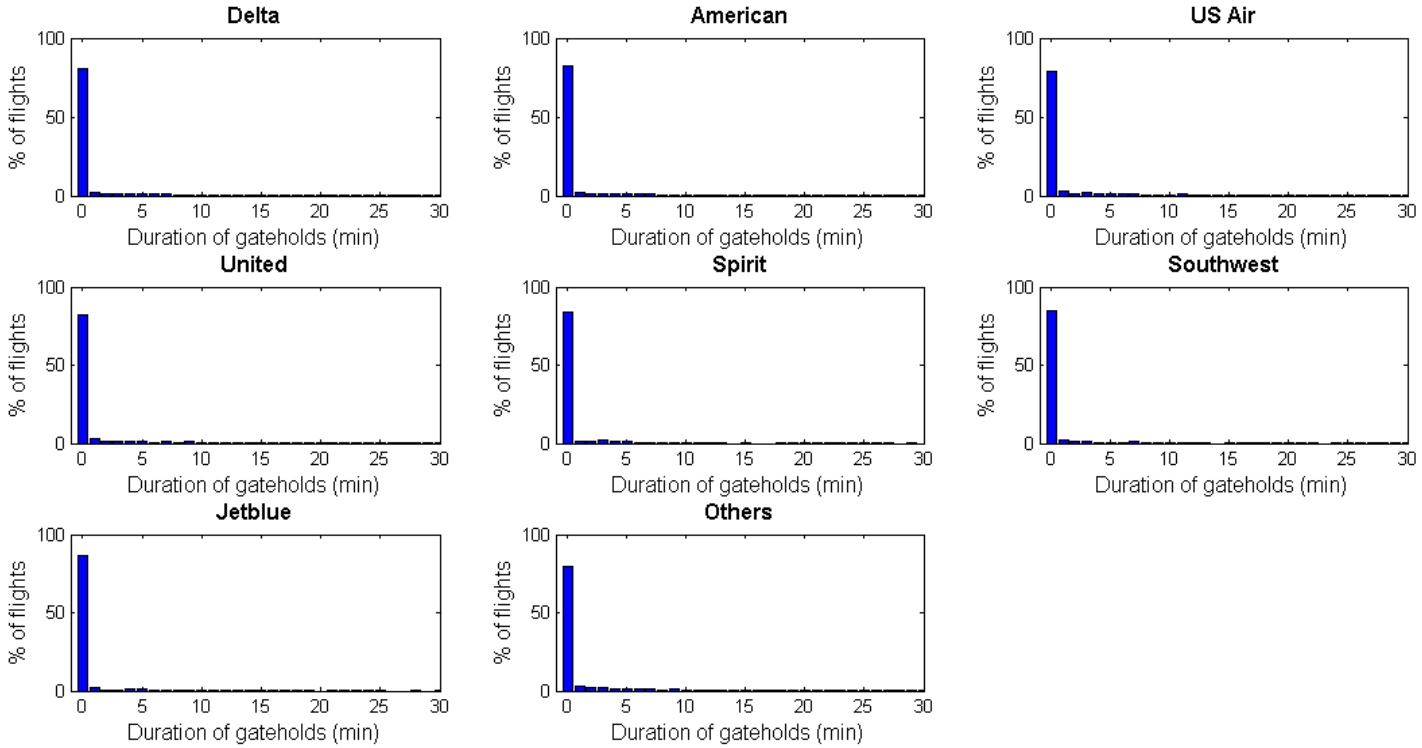

Figure 19: Histogram of duration of gate holds in minutes for the main seven carriers at LaGuardia. Data from July 1, 2013 to August 30, 2013

As depicted in Figure 20, in order to input these restrictions into the model, it is necessary to create a new step in the simulation process. That is, after carrying out the simulation including the pure metering strategy, there needs to be a second simulation that releases those aircraft holding longer than the threshold. Then we can and finally carry out the gate conflict analysis as well as the correction for gate conflicts.

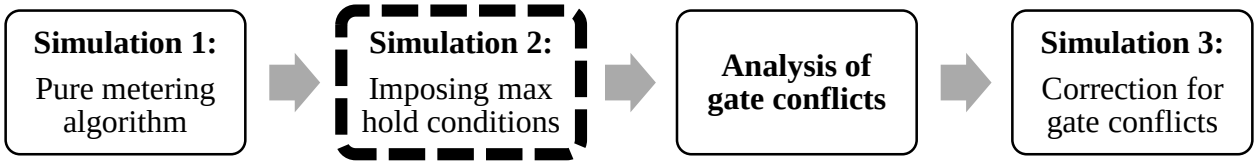

Figure 20: New simulation process to include max hold conditions

At this point, it is relevant to understand how the model goes about implementing this max holding time condition; this can be seen in Table 7. The table shows the new row that would fall between Simulation 1 (pure metering) and the analysis of gate conflicts.

Table 7: Additional row in Table 6 in order to include the maximum holding time condition.

Process	Steps
Simulation 2: Imposing maximum hold conditions	**Step 1:** From the list provided by the gate conflicts analysis, identify all the flights that are affected. **Step 2:** For those flights, change the tobt(i) to the value shown in the second column of the gate conflicts matrix. **Step 3:** For all flights repeat Steps 4, 5 and 6 from Strategy A to allocate the new dqet(i) and atot(i).

The reason for calculating the maximum holding times before the gate conflicts is just a matter of computational cost. Limiting the time aircraft hold at the gate reduces the time aircraft are exposed to cause a gate conflict. Thus, by implementing the MHP policies we reduce the number of gate conflicts generated and, hence, the computational cost of solving those conflicts.

Looking back at all the sections in Chapter 2 it is worth stating that this chapter proposes a metering strategy that tackles the inefficiencies related to queuing time (as opposed to travelling) derived from the pushback at discretion strategy. In particular, metering is based on an accurate understanding of the flow propagations through the airport surface to recommend gate holds, which leads to a reduction in engines-on time during the taxi-out procedure, which in turn leads to reductions in the fuel burn and gas emissions. As main challenges, we can point out the increase in gate conflicts (although Chapter 3 shows the small scale of the increase) and more granular resource management operational challenges that arise from an increase resource utilization time. Early results show that the necessary length to implement such a strategy is in the great majority of cases short; however, a small proportion of flights being held longer times, potentially stressing the challenges of the model, from an airline's perspective. Considering this argument, we introduce three different metering-based policies that limit the length of the gate holds with the rationale that shorter gate holds lead to a significantly less noticeable consequences. The first policy is pure metering, with no gate-holding limit, the second sets a limit of 15 minutes, and the third policy a limit of 10 minutes. With these three policies we believe stakeholders have enough flexibility to choose the option they believe best solves the trade-off between benefits and their own set of challenges. The three strategies present benefits in terms of real money savings (fuel burn reduction), however, the portfolio of strategies allow decision makers to pick the strategy based on the trade-off between the value added and perceived risks. To help inform stakeholders on how

to assess the advantages and disadvantages of implementing a particular policy, Chapter 3 presents numeric values of value added by each policy, in terms of benefits and challenges.

3. RESULTS AND DISCUSSION

This chapter presents the results in the following manner: The first section is a general proof that metering is an effective strategy at reducing the taxi-out time, looking at it from an aggregate perspective, from terminal standpoint, as well as from a runway configuration viewpoint. This first section emphasizes not a quantitative but a qualitative understanding of the taxi-out reduction. The second section provides more detailed analysis about the presence of gate conflicts with the implementation of metering. Finally, the third section presents a more quantitative-based approach to compare the three different maximum gate-holding time policies introduced. It is worth noting that given the way the B.Unr. policy has been defined, it corresponds to the pure metering strategy and thus, the general results in section 3.1 correspond to the unrestricted maximum gate-holding limit policy.

3.1.EFFECTIVENESS OF THE METERING STRATEGY

This section aims at understanding the circumstances in which metering provides a reduction in taxi-out time and thus a reduction in emissions and fuel burn.

3.1.1. GENERAL EFFECTIVENESS

The first aspect to consider is whether the metering strategy actually leads to a significant reduction in taxi-out reduction. In this regard, it is interesting to look at Figure 21, which plots the average taxi-out time by time of the day for both the baseline and metering strategies. These average taxi-out times are computed over the 62 days from July and August 2013. The most interesting aspect to look at in Figure 21 is the difference between the two curves and the relation between this difference and the baseline taxi-out time. The first observation is that once the traffic builds in the

morning (from 6 AM to 9 AM), there is a consistent reduction of taxi-out time throughout the day until the traffic is reduced at night. The reduction is particularly salient during the afternoon peak (4 PM – 6PM). It is worth understanding that these results are average reduction over a period of 62 days, and therefore, the results do not imply that every day will experience a taxi-out a taxi-out reduction as the one displayed in Figure 21. Indeed, there may be days with no reduction at all, and other days with a larger reduction. The main take-away from this figure is the consistency in the taxi-out reduction throughout the day.

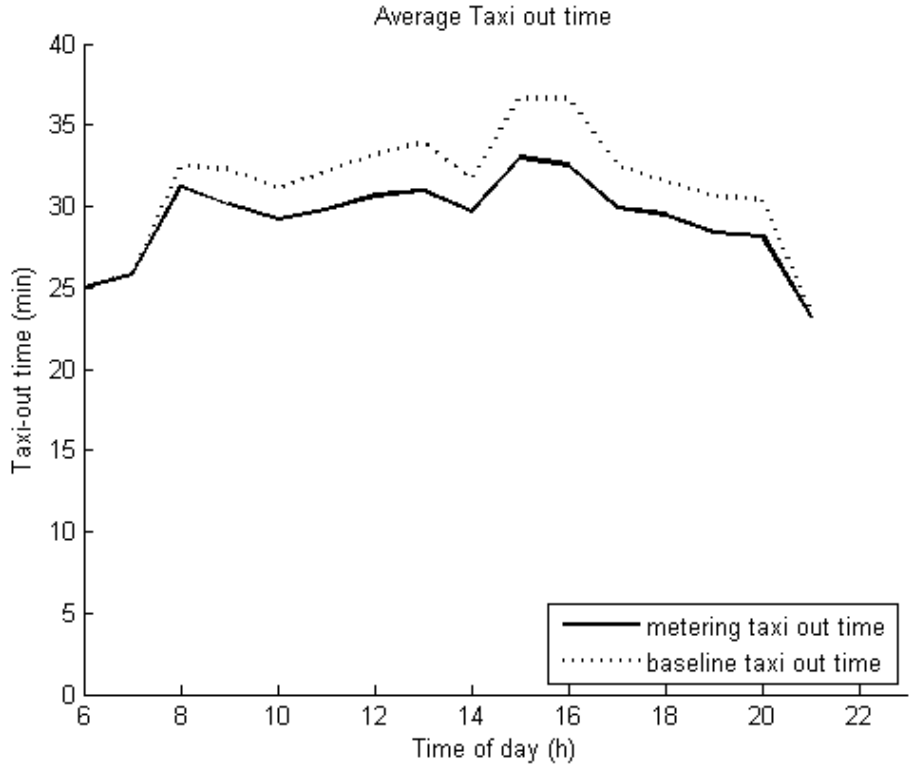

Figure 21: Average taxi-out time by time of the day for the baseline and metering case. The metering strategy leads to a significant reduction in taxi-out time. Data from July 1, 2013 to August 30, 2013

3.1.2. BY TERMINAL

In this section the main goal is to determine the way these taxi-out reductions are distributed among terminals. Figure 22 shows a breakdown of the aggregate data presented in Figure 21 by terminal area. The first observation, and the main take-away, is that flights operating in all terminals are seeing reductions in their taxi-out times throughout the day in a consistent way. Put another way, the reductions seen at an aggregate level are not dependent on the terminal the flight is operated from. In addition it is possible to see that independently from the scheduling levels (number of flights per hour), the flights still experience consistent reductions during the day. This comes from the fact that the bottleneck in this situation is the runway; given that pushback slots and runway take-off slots are allocated on a FCFS basis, the terminal from which the flight is being operated is likely to not be a relevant factor in the taxi-out reduction.

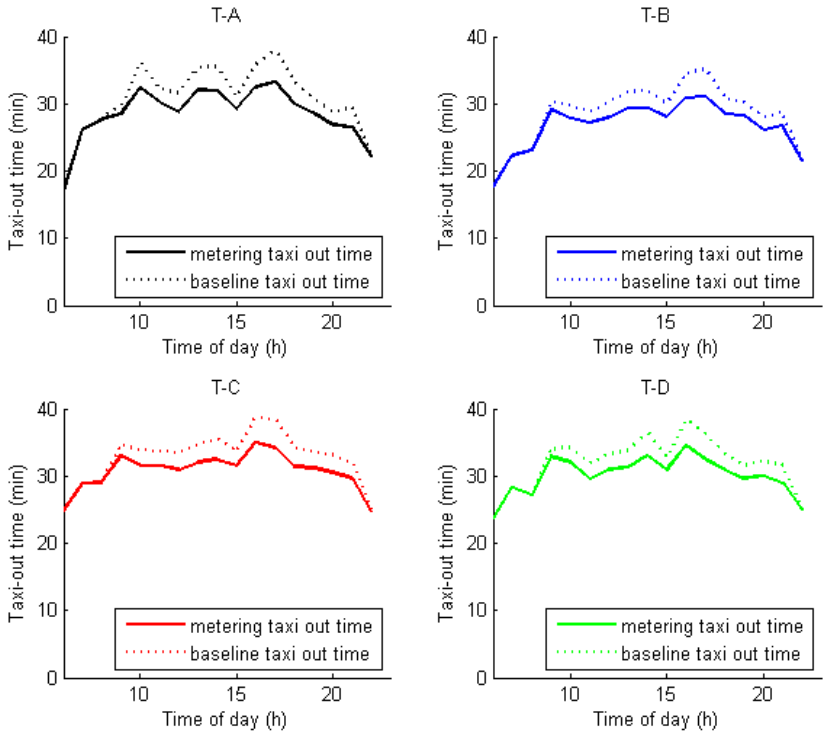

Figure 22: Average taxi-out time by terminal and by time of the day for the baseline and metering case. The metering strategy leads to a significant reduction in taxi-out time. Data from July 1, 2013 to August 30, 2013

3.1.3. RUNWAY CONFIGURATION

The next aspect worth analyzing is the relevance of taxi-out time reductions for the most used runway configurations. Figure 23 depicts a breakdown of the general taxi-out reductions in the five most frequent runway configurations: 31|4, 22|31, 31|31, 22|13, and 4|13. Opposed to the fairly homogeneous reduction in taxi-out time by terminal observed in Figure 22, Figure 23 shows two different behaviors in runway configurations. On the one hand, runway configurations 31|4, 22|13, and 4|13 see fairly consistent taxi-out reductions throughout the day from the morning traffic build-up to the time traffic goes down in the evening. On the other hand, runway configurations 22|31 and 4|13 experience non-existent taxi-out reductions as both the baseline and the metering average taxi-out curves overlap.

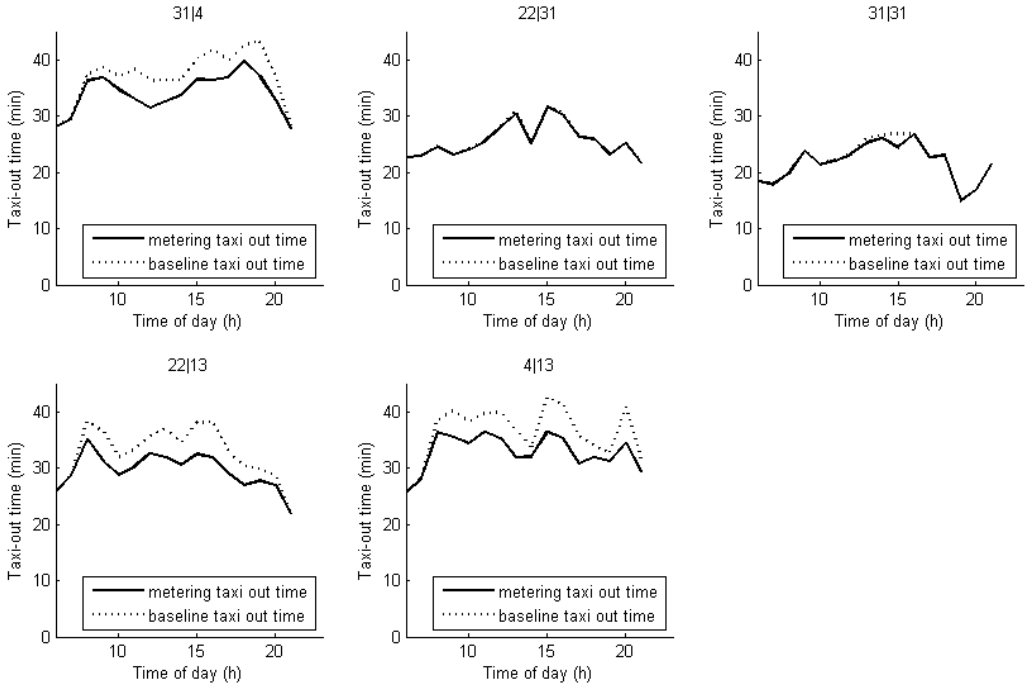

Figure 23: Average taxi-out time by runway configuration and by time of the day for the baseline and metering case. The metering strategy leads to a significant reduction in taxi-out time. Data from July 1, 2013 to August 30, 2013

When comparing the absolute values of baseline taxi-out times (dotted lines) from each subplot in Figure 23, the average taxi-out time values for configurations 22|31 and 31|31 are already lower than the other configurations. This indicates that these runways are being used in circumstances where the surface traffic is low, and thus, congestion is not a salient factor. Given that metering is particularly effective in periods with congestion, we claim that the lack of congestion likely justifies the non-existent or small reduction in taxi-out time.

The objective at this point is to use the data available, ASPM, to justify the difference in taxi-out reduction. We do this by comparing the runway configuration use in the period of analysis, July-August 2013, with those periods of time when metering occurs. Alternatively, the objective is to understand whether the use of all the runway configurations is the same with or without metering. Put another way, the goal is to see whether metering promotes the use of a particular set of runway configurations. To do so, it is worth comparing Figure 24 and Figure 25. On the one hand, Figure 25 depicts the configurations during the July – August 2013 period in which all flights operated; that is, Figure 25 displays the proportion of flights that operated in a particular runway configuration. On the other hand, Figure 24 depicts the configurations during the same period, but only for flights that experienced a gate hold, and thus operated in a period when congestion led to the implementation of the metering strategy. By comparing the two figures, the following conclusions can be drawn:

- Runway configurations 31|4, 22|13, and 4|13 are overrepresented in those periods with metering. These configurations represent more than 90% of the flights with metering, compared to the approximately 70% of those flight in general situations.

- Runway configurations 22|31 and 31|31 are underrepresented during metering as these two configurations are only used approximately 10% of the time, compared to approximately 30% of the time in the general case.

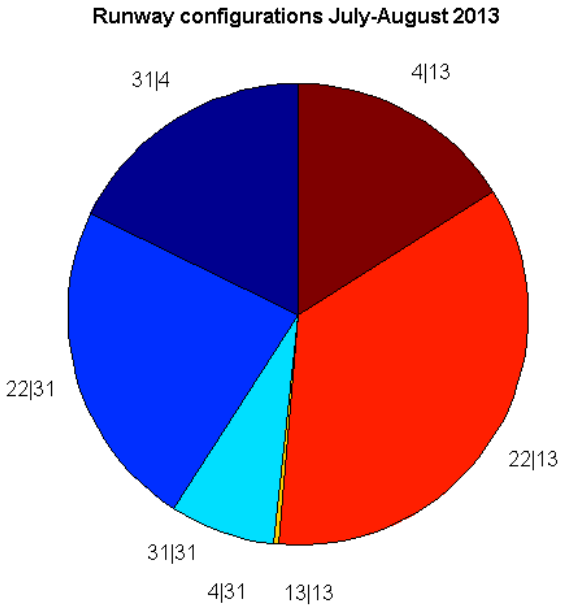

Figure 24: Use of runway configurations during July 2013 and August 2013.Runway configurations 22|31 and 4|31 have approximately a joint 30% share.

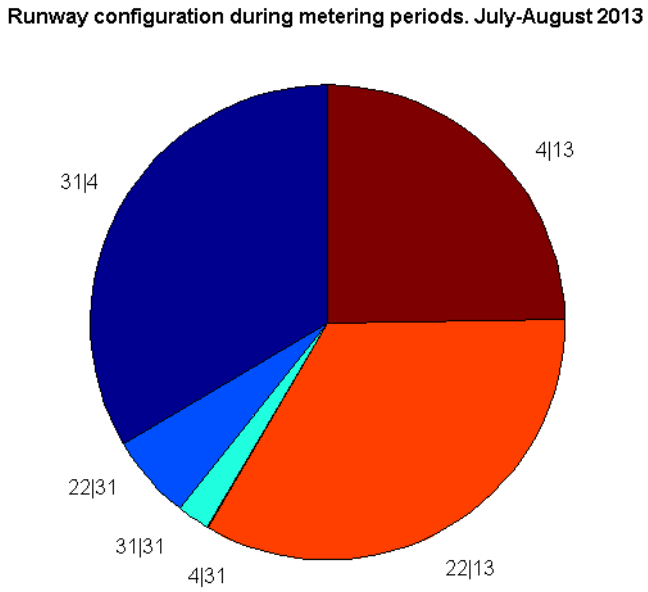

Figure 25: Use of runway configurations during the metering periods on July 2013 and August 2013. Runway configurations 22|31 and 4|31 have approximately a joint 30% share.

From these results we can deduce that there is no significant reduction in taxi-out time because runway configurations 22|31 and 31|31 are barely used when metering and the observed taxi-out times are the shortest, which indicate periods of low congestion. The interesting fact is that both configurations use the same departing runway, runway 31. Therefore, further analysis on runway 31 may suggest the reasons for its not being used much during metering periods. However, this would require additional work beyond the scope of this research.

At this point of it is worth pointing out that metering is an effective strategy to reduce the taxi-out time in a consistent way during the day regardless of the terminal the flights operate from. As for runway configurations, only those configurations used during metering periods see a consistent reduction throughout the day, which also strengthens the effectiveness of the metering strategy.

3.2.GATE CONFLICTS

The previous section proves that metering is effective at reducing the taxi-out time, but it also important to quantify the consequences of the implementation of such a strategy. In particular, this section looks at gate conflicts, which have been described in section 2.4. To this end, Figure 26 shows an airline-specific analysis of average gate conflicts by time of day.

The blue bars show the current gate conflicts obtained from the ASPM and flighstats datasets. That is, the model computes for each gate, the arrivals whose gates are in use at the wheels-on time. In this regard, Delta currently has the worst performance with maximums of 1.3 gate conflicts per hour at 10 AM. In this case, having a larger market share and operations at the airport does not justify the larger number of gate conflicts as Delta also has more gates. The larger number of conflicts has to with the way Delta uses these gates and allocates flights to them. The other carriers

currently present a significantly lower level of conflicts around 0.4 gate conflicts per hour, on average.

As expected, holding aircraft at the gate leads to additional gate conflicts, which are depicted in red in Figure 26. Although the initial expectations were to see a significant increase in gate conflicts, the actual numbers are small. In particular for Delta, the maximum average additional number of gate conflicts is 0.4 conflicts per hour, a situation that only concerns one one-hour period during the day; the rest of the day, the values are significantly smaller. An average value of 0.4 conflict is small because out of the approximately 40 departures per hour, only 0.4 is affected by this type of conflicts. For the case of Delta, there is a sustained 0.1 to 0.4 additional gate conflicts throughout the day, particularly from the moment the traffic builds up in the morning to the moment the traffic decreases in the evening. This is an interesting case as the gate conflicts (direct cost) are sustained throughout the day, and the same occurs with taxi-out reductions (direct benefit). Such a behavior does not replicate for other carriers. On the one hand, American and US Airways have two hours (3 PM to 5 PM) when the gate conflicts are on the 0.2 to 0.4 level, but the rest of the day the gate conflicts are disregardable. As for the rest of the carriers, the additional gate conflicts are either non-existent or they are too small to have a considerable impact on the performance of these carriers. Such low results are the consequence of the comparatively low scheduling levels at the airport out of gates where they are the only carriers, which allows them to schedule a small number of flights in a more robust manner.

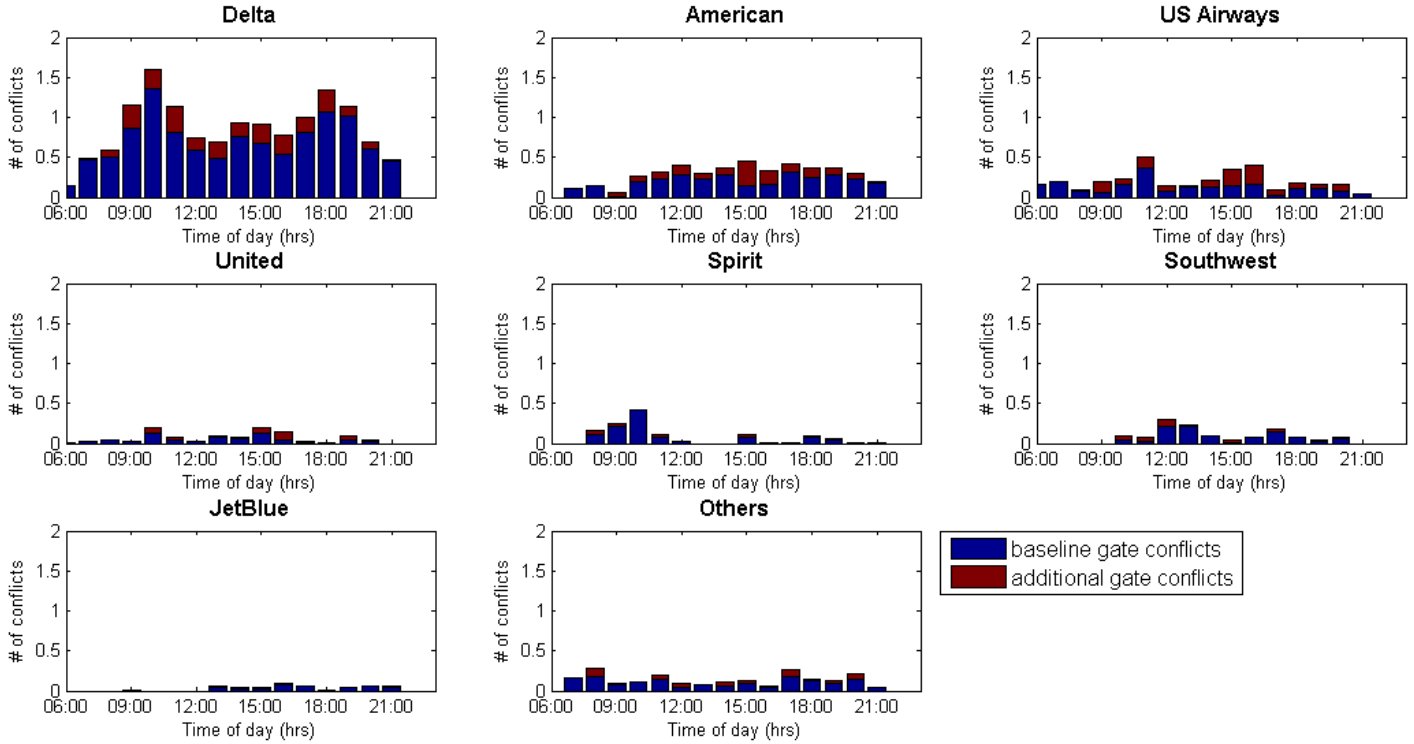

Figure 26:Baseline (Strategy A) gate conflicts by time of day in blue and additional gate conflicts by time of day when implementing Metering (Strategy B.Unr.). Data from July 1, 2013 to August 30, 2013

After understanding the small scale impact of the gate conflicts derived from metering, it is worth understanding whether there is any particular variable that is more influential when analyzing gate conflicts. Indeed, the total number of gate conflicts is a function of many variables such as the scheduling strategy, the limited number of gates, the turnaround times and the network delays. When it comes to understanding gate conflicts, other values may come into play. Our initial hypothesis are to analyze the time of day and the weather as potential sources of gate conflicts. The results from Figure 27 and Figure 28, however, do not support our initial assumptions.

First, Figure 27 illustrates the average number of gate conflicts by time of day (in black) and the same indicator broken down by weekday and weekend. As the black curve is an average value, it is interesting lo break it down in subsets to identify trends based on the subset. In particular Figure 27 shows no clear evidence as to whether more gate conflicts occur during weekdays or weekends with the only potential exception of Delta, which shows a smaller level of conflicts on weekends than weekdays in some periods of the day, although not throughout the day. The potential justification for such a difference is the regular 30% scheduling reduction that Delta implements on Saturdays to accommodate the reduction in demand. As indicated earlier, the evidence of the difference in pattern between weekdays and weekends is not strong and even in regions where the difference is strong, the smaller value of gate conflicts is negligible.

The rest of the carriers present a similar, or not discernable, trend between the two sorts of days.

Second, Figure 28 depicts the average number of gate conflicts by time of day (in black) and the same indicator broken down by two weather categories, whether or not there were RAPT values larger than zero. Put another, “days with RAPT” category contains all the days with some RAPT values above 0. To reiterate, this research uses the RAPT variable as a proxy for weather. As for the results, Figure 28 shows that there is not enough evidence to think of significant difference

between days with or without positive RAPT values. Similarly to what happened with the day of the week factor, only Delta seems what could be regarded as a difference between periods with weather and without weather, but the evidence is weaker than in the previous case; indeed, the difference is only observable in the afternoon and evening times.

Based on this analysis we can state that metering leads to additional gate conflicts that are smaller than initially expected and in average values per hour, they are small when compared to traffic data. In trying to identify factors influencing the presence of gate conflicts, the day of the week, and the weather, the data reveals no evidence to help us divide the data into segments in order to better predict behavior.

This extensive analysis shows the low frequency of gate conflicts, which allows us to manage these events without significantly altering airport surface operations nor diminishing the effectiveness of the proposed strategy.

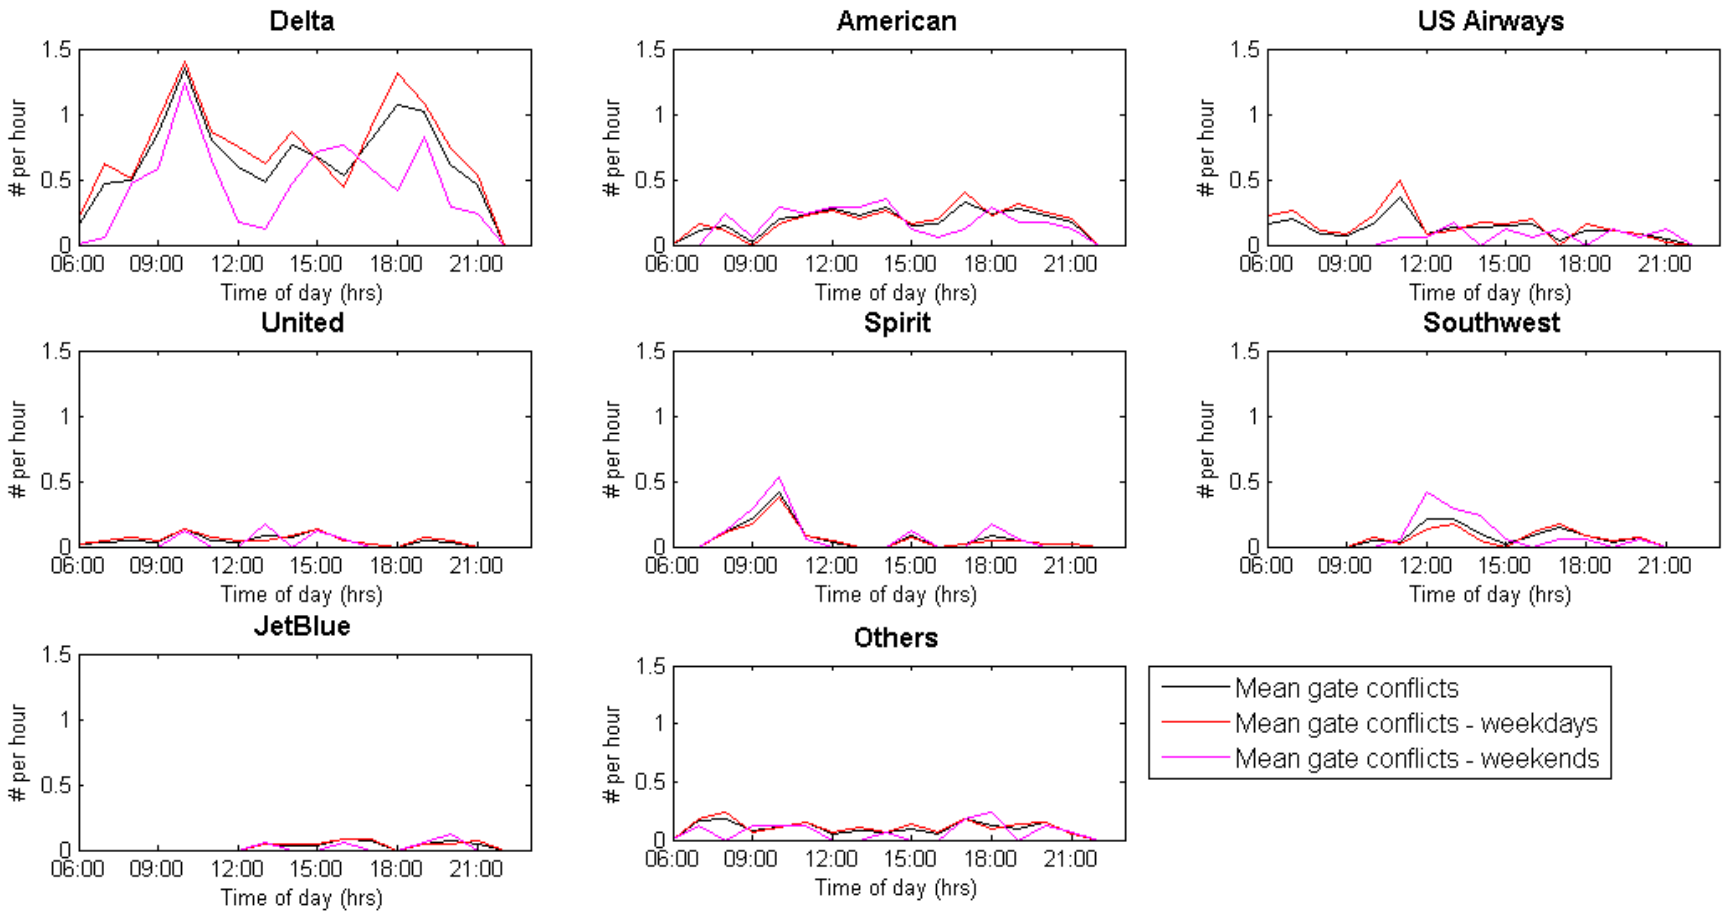

Figure 27: Mean gate conflicts by time of day in three situations: mean of all days, week days and weekends. Data from July 1, 2013 to August 30, 2013

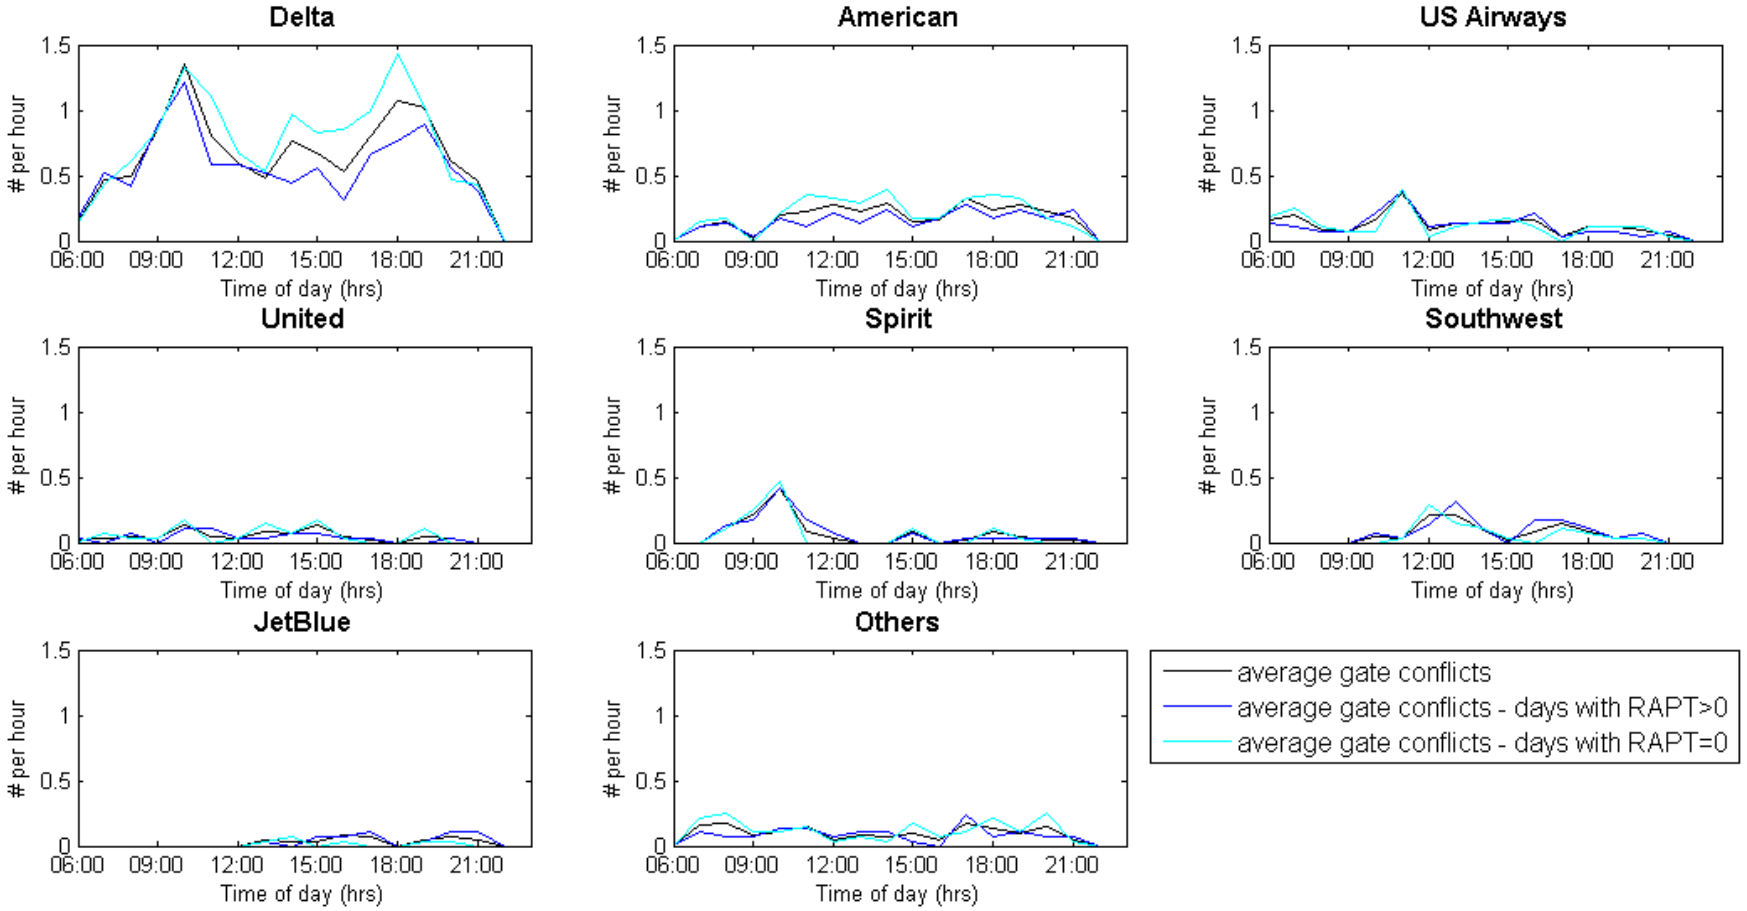

Figure 28: Mean gate conflicts by time of day in three situations: mean of days with RAPT larger than 0 (“bad weather”), mean of days with RAPT equal to 0 (“good weather”). Data from July 1, 2013 to August 30, 2013

3.3.ANALYSIS OF MAXIMUM GATE-HOLDING POLICIES

After understanding the dynamics of the taxi-out reduction and the occurrence of gate conflicts, this section analyzes quantitatively the implementation of the three proposed metering-based maximum holding policies. Going back to the definition of these maximum holding policies, this analysis compares the B.Unr. or unrestricted policy, the B.15, and the B.10 policies. The first does not limit the gate-holding time at all and therefore the results shown for such a policy are those of the pure metering strategy. The second and third policies limit the gate-holding times to 15 and 10 minutes, respectively. As described earlier, these three proposed strategies aim at providing the stakeholders with flexibility around the implementation process. With these policies on the table, each stakeholder can understand how those interact with internal operational practices and weigh the benefits and costs.

Before delving into the results in Figure 29, it is worth taking into consideration the following observations. First, the color code used is: black corresponds to the baseline case or Strategy A (push-back-at-discretion), blue corresponds to Policy B.Unr, red corresponds to Policy B.15, and finally green corresponds to Policy B.10. Second, there are three sets of curves plotting three different variables:

- The topmost subplot shows surface traffic for the four curves, with surface traffic defined as the number of aircraft that, at the beginning of the time period, have started the pushback procedure (not including those that are holding at the gate) but have not yet taken off. From a mathematical perspective, the surface traffic at time period T is the sum of all flights that fulfil the two following conditions:
    - the TOBT(i), the Target Off-Block Time, is earlier (or smaller) than the initial time t of time period T

    - the ATOT(i), the Actual Take-Off Time, is later (or larger) than the initial time t of time period T

- The curves on the positive side of the lower subplot depict average taxi-out times, defined as the difference between atot(i) and tobt(i).
- The curves on the negative side of the lower subplot illustrate average gate-holding times. Showing these gate-holding times in the negative area of the subplot is a way to easily plot the three variables together, as they are clearly interwoven.

Despite plotting the three variables together, it is important to be aware of the time lapse between them. The simulation time period length of 15 minutes, and thus, if an aircraft with EOBT is T is held during T, it is likely that this aircraft does not start taxiing-out until time period T+1 and therefore the reduction in surface traffic is not observed until period T+1 or even T+2. That is the reason why some accounting mismatches of one or two time periods occur.

At this point, it is worth understanding how each policy works by analyzing Figure 29, in particular, the connections between the three variables displayed. The main observations from Figure 29, corresponding to July 9, 2013, are the following:

- For surface traffic levels below the N_{ctrl} level (21 in the case of the Figure 29), the three metering policies yield the same results as those provided by the baseline case. This translates into the color curves overlapping with the black curve. Such behavior can be seen from the beginning of the day until 11:45AM.
- The three metering policies behave similarly in circumstances when the surface traffic is slightly above of the N_{ctrl} level. These small differences lead to short gate holds, which are shorter than 10 minutes, and thus, policy B.10 does not require any forced push. This situation takes place between noon and 2:30 PM. Despite the small size of the difference

in the figure, the real magnitude of the taxi-out reductions, which is similar to that of the gate holds, is about 5 minutes.

- In periods when the surface traffic is significantly above the N_{ctrl} value, there is a significant difference between the unrestricted policy and the other two policies. The three policies perform the same up to a point at which the B.10 policy starts noticing aircraft being held more than 10 minutes and forces those aircraft to pushback even though the metering strategy would recommend them to hold at the gate, given that they are generating congestion. A similar argument can be made for the B.15 policy when aircraft start holding longer than 15 minutes.
- B.10 and B.15 are effective at maintaining gate-holding times at 10 and 15 minutes respectively. The flat lines green and red lines during the 3:30 PM to 8:30 PM are clear signs of the model's effectiveness of ensuring there is no gate hold longer than the specific threshold.
- The unrestricted policy is the most effective at keeping the surface traffic constant throughout the day at values around the N_{ctrl} level. However, the performance during the evening peak is above the expected threshold; the reason for such higher-than-expected traffic levels is the presence of gate conflicts as gate holds increase. As explained in section 2.4, gate conflicts are resolved, forcing a pushback regardless of whether this additional aircraft on the surface brings the surface traffic to exceed N_{ctrl}.
- Having taxi-out values and gate holds on the same scale allows for fair comparisons between reductions and gate holds. Before looking at quantitative results (analyzed below), it possible to state that taxi-out reductions seem comparable to the gate-holding times, plus

or minus some adjustment of one or two time periods due to surface traffic taking place one or two time periods after the EOBT of each flight.

- In periods with large surface traffic, policies B.10 and B.15 perform similarly; specifically, they lead to parallel taxi-out and gate-holding curves separated 5 minutes from one another.

The space between the black and blue curves in the three families of curves show the reduction opportunities by implementing different maximum gate-holding strategies. Indeed, this research has chosen 10- and 15-min thresholds, but should other values be picked, the resulting performance curves would fall somewhere between the black and blue curves. However, due to the non-linearity of the problem it is difficult to pick a gate-holding limit to obtain particular threshold.

These observations are made based on the data from July 9, 2013, which was chosen because of the clarity of the values and the figures associated to it. However, not all the days in the data behave in such a way and actually, there are days in which there are no gate holds at all because the surface traffic is never above the N_{ctrl} level. As a reminder, the high values of surface traffic may be due to two main reasons:

- The number of aircraft scheduled to take-off in a particular tie period is large.
- The weather and or airport characteristics in that particular period lead to an airport capacity which is unable to manage given demand.

The former reason is a given or static reason, as scheduling takes place at least six months in advance, but the latter reason is dynamic as it depends on how is the weather like on a particular day or what runway configuration the airport is operating in at that particular moment. Therefore, one particular schedule may be operated under no congestion or gate holds, and that same schedule may lead to significant congestion and gate holds given that the dynamic part of the problem is being affected by weather. However, in reality these two problems are coupled as although

scheduling levels tend to remain fairly consistent throughout a season, the real performance of that schedule can vary significantly from what is predicted due to weather, and runway configuration to a lesser extent. Indeed, weather affects the overall airport capacity, including arrivals and departures. Such modifications in capacity are likely to change the realization of the schedule from what they are initially planned. To this end, the evening surface traffic levels observed in Figure 29 are a consequence of the decrease in capacity due to RAPT values higher than 0 that lead to lower departure throughput capacity.

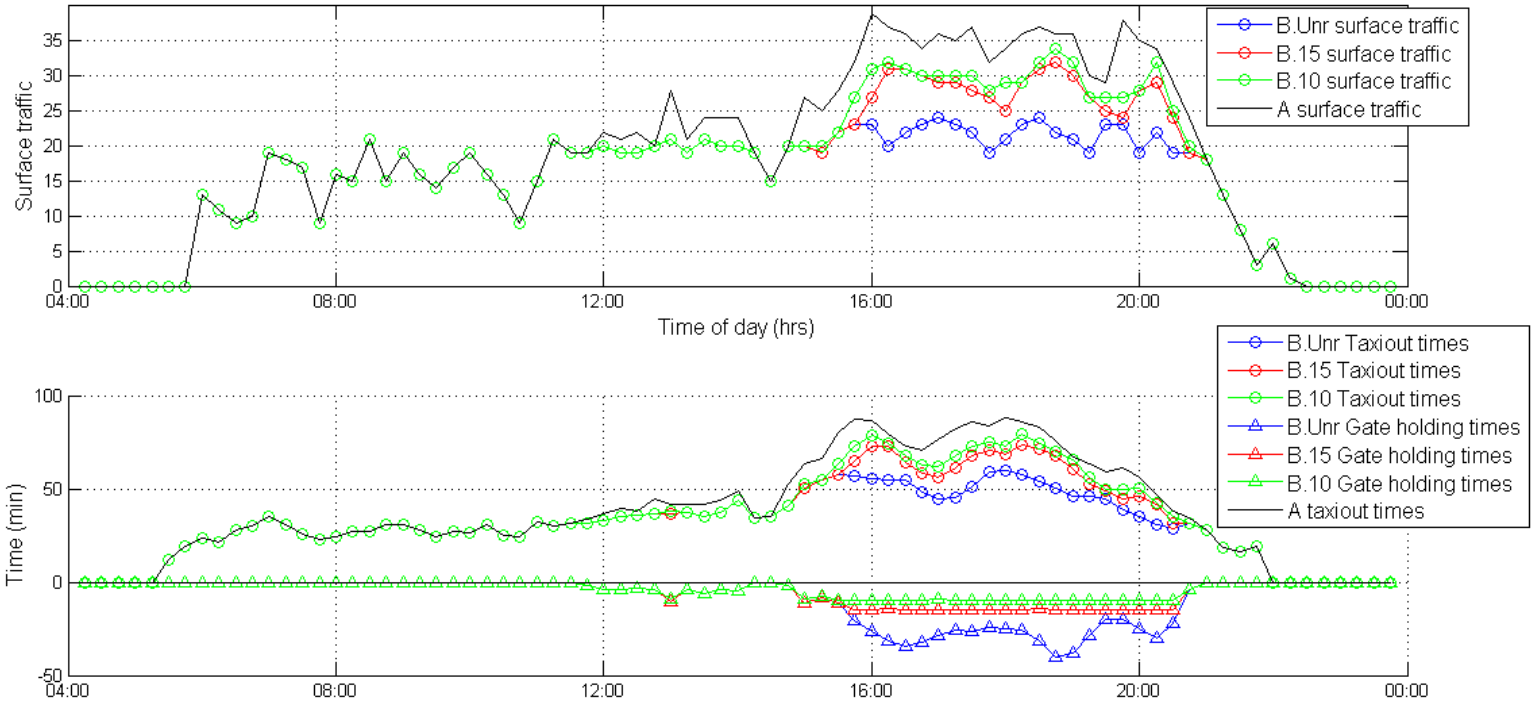

Figure 29: Surface traffic, taxi-out times and gate-holding times for three MHP, compared with the baseline cases. Data from July 9, 2013

Having analyzed Figure 29 and conceptually derived the dynamics behind the different policies, Table 8, Table 9, and Table 10 quantify the benefits, costs and fairness of each policy in terms of time and benefit share. Principally, the information provided in the three tables is organized as follows:

- Taxi-out time reductions: In this block we provide the data to assess the benefits of each metering policy.
    - Total airline taxi-out reduction over the 62-day period (in minutes): This value is the addition of all the taxi-out reductions by airline for all flights from July 1 to August 31 2013 (i.e. Delta is expected to see a reduction of 24,493 minutes for all the Delta flight belonging to the July-August period should the unrestricted policy be implemented). From the model's performance viewpoint, this number itself does not provide much information, apart from showing how large the number can become over the period of two months. However, from the operator's perspective, this is number can be multiplied by unit emissions and unit fuel burn to obtain total reductions.
    - % Reduction in taxi-out time: This is % reduction in total taxi-out time when compared with the baseline case for each particular airline. (i.e. For Delta, the 24,493 minutes would represent a 6.5% reduction should the unrestricted policy be implemented). This number is relevant as it is one of the key performance indicator of the benefits associated with a particular policy.
    - Percentage of the benefits: This value is the share of the taxi-out reductions that corresponds to each airline in percentage terms. That is, if all the taxi-out reductions, in minutes, for each airline are added up we would obtain all the benefits

generated by each policy; the percentage of the benefits is the portion of the total benefits allocated to every airline (i.e. For Delta, the predicted 24,493 reduced minutes if the unrestricted policy were implemented would represent a 39.1% of the benefits achieved by all the carriers). This indicator is relevant while assessing fairness of benefits; further information is provided in the Market Share block below.

- Gate Holds: In this block we provide data to evaluate the impact of each of metering policy in terms of the gate holds generated, in absolute and relative terms.
    - Total gate holds over the 62-day period (in minutes): This indicator adds up all the gate-holding times that all the flights for each airline experience over the course of the 62 days (i.e. The unrestricted policy is expected to lead to 24,660 minutes of holding for Delta). From the model's performance viewpoint, this number itself does not provide much information, apart from showing how large the number can become over the period of two months.
    - % of the benefits: This variable computes the airline-specific portion of the total gate-holding time (i.e. The model predicts that Delta would experience 39.4% of the total gate-holding time by all airlines). This indicator is relevant while assessing fairness of the gate holds; further information is provided in the Market Share block below.
- Market Share of departures: This indicator accounts for the share of departures operated by each airline during the July-August 2013 period (i.e. Delta operated 38% of the departures). This indicator is the base for comparing the percentage of benefits and

percentage of the gate holds that each airline experiences, which will allow us to make conclusions about the fairness of the metering strategy.

Understanding the variables, it is possible to start learning from each table. First, Table 8 displays the performance data for the unrestricted policy, which is equivalent to a pure metering strategy and can be seen as the maximum yield that can be derived from implementing metering. In such conditions, the main key points are:

- The average efficiency of this policy is between 5.7% and 6.9%. Interestingly, the values on the lower side of the interval belong to smaller carriers and those in the higher part belong to bigger carriers.
- The unrestricted policy leads to a reduction of taxi-out times by airline that are commensurate with the gate-holding time.
- The percentage of benefits are nearly identical to the percentage of the costs, which is a clear signal of fairness. That can be thought of as, for each 1% of the total taxi-out reduction an airline needs to wait 1% of the total gate-holding time. From the previous point, we know that these two total values are the same and thus, from 1 minute of holding airlines get 1 minute of taxi out reductions.
- The percentage of benefits (and thus the percentage of the gate holds) for each airline is approximately equal to the share of departures. For example, Delta operated 38% of the departures in the period under study and the model grants Delta 39.1% of the benefits of holding and 39.4% of the gate holds. The error between the benefits and the departure share is +/- 1.5% which we believe is a small prediction error. Based on this finding, we can state that this policy ensures fairness, given that operators with more departures are likely to see more benefits (and more gate holds) from the model.

In addition to these observations, it is useful to keep in mind what are the main consequences of this policy in terms of gate holds, particularly the distribution of their duration. For the unrestricted policy, or pure metering, the distribution of gate hold duration is shown in Figure 19, showing the preliminary results of the pure metering strategy. Having that distribution in mind is helpful to understand the performance of the two policies presented below.

The B.15 policy performance results are shown in Table 9. Most of the observations from in this case are rather similar to those from the unrestricted policy, except for the magnitude of the total taxi-out reductions and efficiency rates that on this occasion fall between 2.8% and 3.8%. The benefits are commensurate with the costs, which in turn are approximately equal to the departure share, ensuring the fairness of the policy.

Keeping in mind the distribution of gate hold durations from Figure 19, it is helpful to see how policy B.15. changes the shape of that distribution; such a change in the shape of the histogram is depicted in Figure 30. The only difference when compared with Figure 19 is that all the values located in the tail to the right of the 15-min threshold are moved to the 15-min bar. Put alternatively, any gate hold longer than 15 minutes is reduced to 15 minutes; that is the reason why the frequency associated with 15 minutes increases. The interesting part is that the left part of the plot remains the exactly the same.

.

Table 8: Taxi-out time reductions (absolute and relative) and gate-holding time (absolute and relative) for policy B.Unr. Data from July 1, 2013 to August 30, 2013

B.Unr	TAXI-OUT TIME REDUCTIONS			GATE HOLDS		Market share
	Minutes	Reduction	% of the benefits	Minutes	% of gate holds	
Delta	24,493	6.5 %	39.1 %	24,660	39.4 %	38.0 %
American	11,713	6.6 %	18.7 %	11,662	18.6 %	19.8 %
US Airways	5,842	6.9 %	9.3 %	5,848	9.3 %	8.2 %
United	2,377	6.4 %	3.8 %	2,343	3.7 %	4.3 %
Spirit	1,037	6.0 %	1.7 %	1,038	1.7 %	2.0 %
Southwest	2,486	5.8 %	4.0 %	2,489	4.0 %	5.1 %
JetBlue	1,450	5.7 %	2.3 %	1,421	2.3 %	3.1 %
Others	13,250	6.9 %	21.1 %	13,187	21.1 %	19.6 %

Table 9: Taxi-out time reductions (absolute and relative) and gate-holding time (absolute and relative) for policy B.15. Data from July 1, 2013 to August 30, 2013

B.15	TAXI-OUT TIME REDUCTIONS			GATE HOLDS		Market share
	Minutes	Reduction	% of the benefits	Minutes	% of gate holds	
Delta	13,345	3.5 %	39.0 %	13,392	39.4 %	38.0 %
American	6,472	3.6 %	18.9 %	6,262	18.4 %	19.8 %
US Airways	3,207	3.8 %	9.4 %	3,212	9.4 %	8.2 %
United	1,267	3.4 %	3.7 %	1,263	3.7 %	4.3 %
Spirit	572	3.3 %	1.7 %	537	1.6 %	2.0 %
Southwest	1,322	3.1 %	3.9 %	1,313	3.9 %	5.1 %
JetBlue	710	2.8 %	2.1 %	699	2.1 %	3.1 %
Others	7,296	3.8 %	21.3 %	7,345	21.6 %	19.6 %

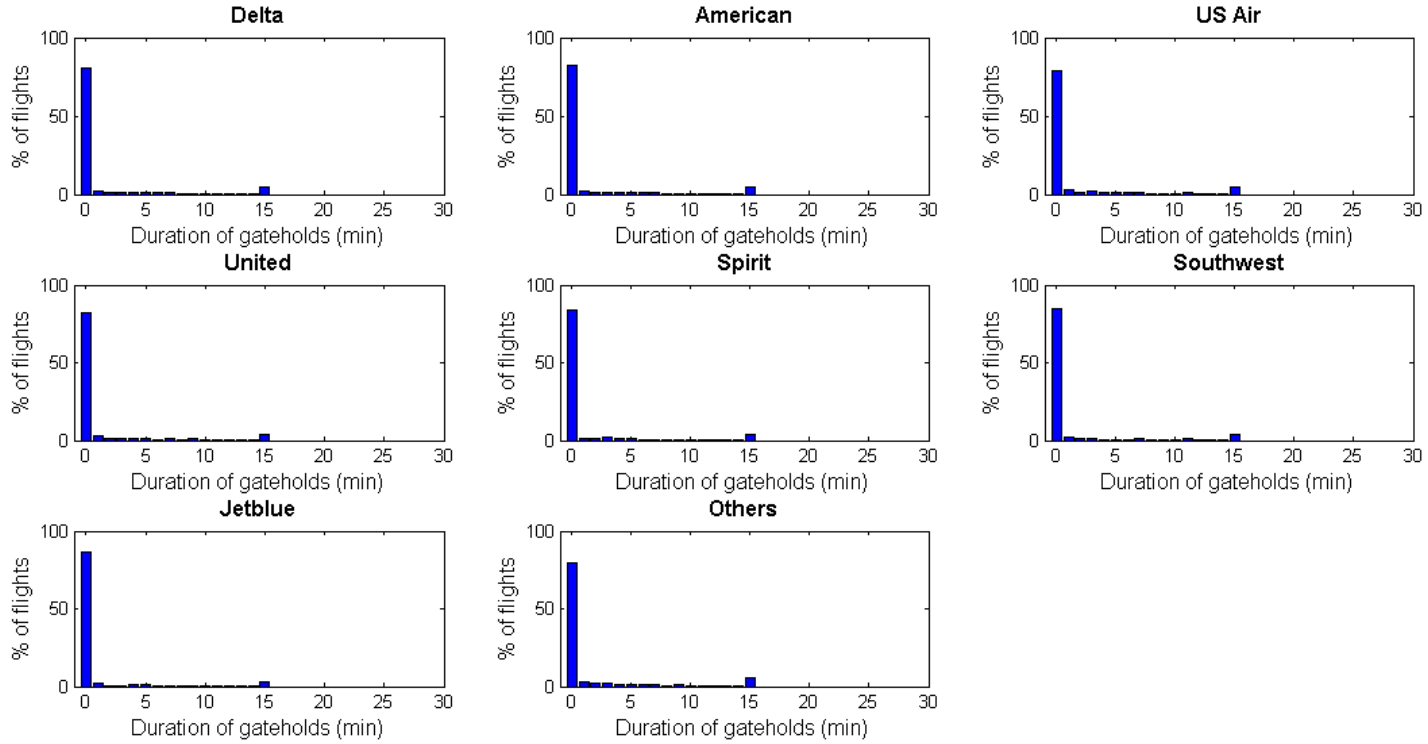

Figure 30: Policy B.15. histogram of duration of gate holds in minutes. Data from July 1, 2013 to August 30, 2013

The simulation results for policy B.10 are shown in Table 10. Two main observations can be made: first, the taxi-out reduction efficiency rates range from 2.4% to 3.1%; second, as happened with policies B.Unr. and B.15, the distribution of benefits and gate holds is approximately the same as the distribution of departure shares, ensuring fairness.

Another interesting aspect to observe is the new distribution of gate-holding times resulting from policy B.10, depicted in Figure 31. As predicted, there are no gate holds longer than 10 minutes, and all the observations that were initially to the right of the 10-minute point are relocated around the 10-minute point.

Table 10: Taxi-out time reductions (absolute and relative) and gate-holding time (absolute and relative) for policy B.10. Data from July 1, 2013 to August 30, 2013

B.10	**TAXI-OUT TIME REDUCTIONS**			**GATE HOLDS**		**Market share**
	Minutes	**Reduction**	**% of the benefits**	**Minutes**	**% of gate holds**	
Delta	11,547	3.1 %	38.9 %	11,602	39.3 %	38.0 %
American	5,601	3.1 %	18.8 %	5,444	18.4 %	19.8 %
US Airways	2,815	3.3 %	9.5 %	2,763	9.4 %	8.2 %
United	1,085	2.9 %	3.7 %	1,103	3.7 %	4.3 %
Spirit	512	3.0 %	1.7 %	478	1.6 %	2.0 %
Southwest	1,139	2.7 %	3.8 %	1,133	3.8 %	5.1 %
JetBlue	606	2.4 %	2.0 %	615	2.1 %	3.1 %
Others	6,409	3.3 %	21.6 %	6,406	21.7 %	19.6 %

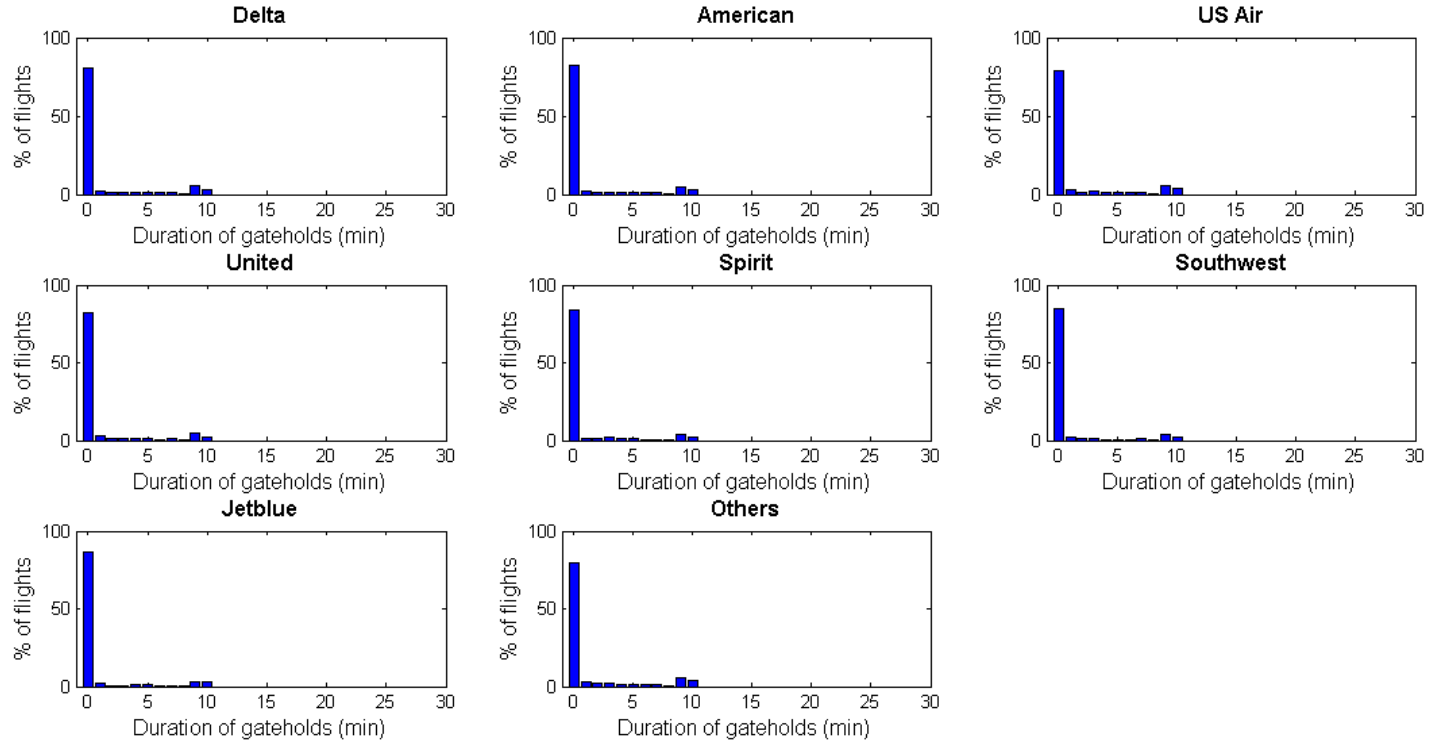

Figure 31: Policy B.10. histogram of duration of gate holds in minutes. Data from July 1, 2013 to August 30, 2013

The natural next step is to compare the three strategies using the indicators introduced in Table 8, Table 9, and Table 10. To this end, Table 11 brings together all this information, rearranging it to present it in a more clear way. The main considerations from Table 11 are the following:

- Limiting the gate-holding time to 15 minutes leads to a reduction of benefits (taxi-out reductions) and costs (gate-holding time) of about 45%; for example, Delta is expected to go from 24,493 minutes of taxi-out reduction to 13,345 minutes, which falls approximately in the 45% reduction. The same calculations can be done for the total gate-holding times. In addition, these calculations hold true for all the carriers.
- Similarly to the previous point, limiting the gate-holding time to 10 minutes, starting from an unrestricted policy, leads to a reduction in benefits and costs of about 55%.
- As expected, all the taxi-out reduction efficiency rates decrease as the policies restrict more; that is, the less time we allow aircraft to hold, the less the taxi-out reduction. The fourth column of the table quantifies that trend.
- The distribution of benefits are split nearly equally (with an error of 0.2%) between airlines regardless of the metering policy implemented. The same is true gate-holding time.
- These distributions of benefits and costs are approximately the same as the distribution of departures among airlines (departure share).

Table 11: Table summarizing Table 8, Table 9, and Table 10. Data from July 1, 2013 to August 30, 2013

Airline	Policy	Taxi-out time reductions			Gate-holding time		% share of departures
		Minutes	% reduction	% share	Minutes	% share	
Delta	Unr.	24,493	6.5	39.1	24,660	39.4	38.0%
	15	13,345	3.5	39.0	13,392	39.4	
	10	11,547	3.1	38.9	11,602	39.3	
American	Unr.	11,713	6.6	18.7	11,662	18.6	19.8%
	15	6,472	3.6	18.9	6,262	18.4	
	10	5,601	3.1	18.8	5,444	18.4	
US Airways	Unr.	5,842	6.9	9.3	5,848	9.3	8.2%
	15	3,207	3.8	9.4	3,212	9.4	
	10	2,815	3.3	9.5	2,763	9.4	
United	Unr.	2,377	6.4	3.8	2,343	3.7	4.3%
	15	1,267	3.4	3.7	1,263	3.7	
	10	1,085	2.9	3.7	1,103	3.7	
Spirit	Unr.	1,037	6.0	1.7	1,038	1.7	2.0%
	15	572	3.3	1.7	537	1.6	
	10	512	3.0	1.7	478	1.6	
Southwest	Unr.	2,486	5.8	4.0	2,489	4.0	5.1%
	15	1,322	3.1	3.9	1,313	3.9	
	10	1,139	2.7	3.8	1,133	3.8	
JetBlue	Unr.	1,450	5.7	2.3	1,421	2.3	3.1%
	15	710	2.8	2.1	699	2.1	
	10	606	2.4	2.0	615	2.1	
Others	Unr.	13,250	6.9	21.1	13,187	21.1	19.6%
	15	7,296	3.8	21.3	7,345	21.6	
	10	6,409	3.3	21.6	6,406	21.7	

Based on the analysis so far, several conclusions can be drawn:

- One minute of gate-hold leads approximately to one minute reduction in the taxi-out time. Such a fact is important as it ensures that no additional departure delay is added to the taxi-out time in the baseline case. Put another way, the gate-holding time is subtracted from the idling time without causing any delay in taxi-out.
- From the three policies analyzed, setting stricter limits leads to different levels of total taxi-out reductions - below we prove that the reduction follows a non-linear trend. However, regardless of the reduction, the benefits and costs of metering are commensurate, and the policies are fair regarding the distribution of benefits and costs with airline departure share.

Based on these considerations, we assume that the one-to-one relationship of cost and benefits, and its distribution fairness remain, independently of the gate-holding limit. With such an assumption we can calculate, without the detailed approach used until now, the reduction for other gate-holding limits different than the ones already calculated, in order to understand how the total taxi-out reduction varies based on the gate-holding limit set. To do so, we run simulations for 5-, 10-, 15-, 20-, 30-, and 35-min thresholds and we compute the average, the minimum and the maximum taxi-out reduction experienced by airlines. Figure 32 plots these taxi-out reductions as a function of the gate-holding limit. Apart from the model performance value, Figure 32 can be extremely useful during the pre-implementation decision making process. Indeed, the table connects the operational effort an airline is willing to make with the benefits associated with that effort. The operational effort consists of ensuring that the airline can manage a gate hold of a particular length, both in terms of having enough handling resources to deal with gate holds and in terms of scheduling through planning gate use by adding some additional buffer time to accommodate those gate holds. The benefits' side is key as it can directly translate into economic

value because percentage reduction in taxi-out can directly translate into fuel burn reductions during taxi-out.

Such considerations allow us to analyze Figure 32 from both the model performance and the decision maker's perspective. The most relevant comments on the figure are:

- The limit-benefit curve presents a non-linear behavior. That is, equal movements in the x-axis lead to different benefits changes in the y-axis; were this relation linear, then the unit change in values would lead to the exact same change in benefit, no matter where in the x-axis the evaluation point is located.
- The average reduction tracks rather accurately the trends set by the minimum and the maximum. Such an observation is particularly salient to assess the reliability of the average as a good indicator.
- Incremental benefit after the 20-min gate-holding limit are marginal. Indeed, the low frequency of gate holds longer than 20 minutes makes longer gate-holding limits to not contribute to much benefit.

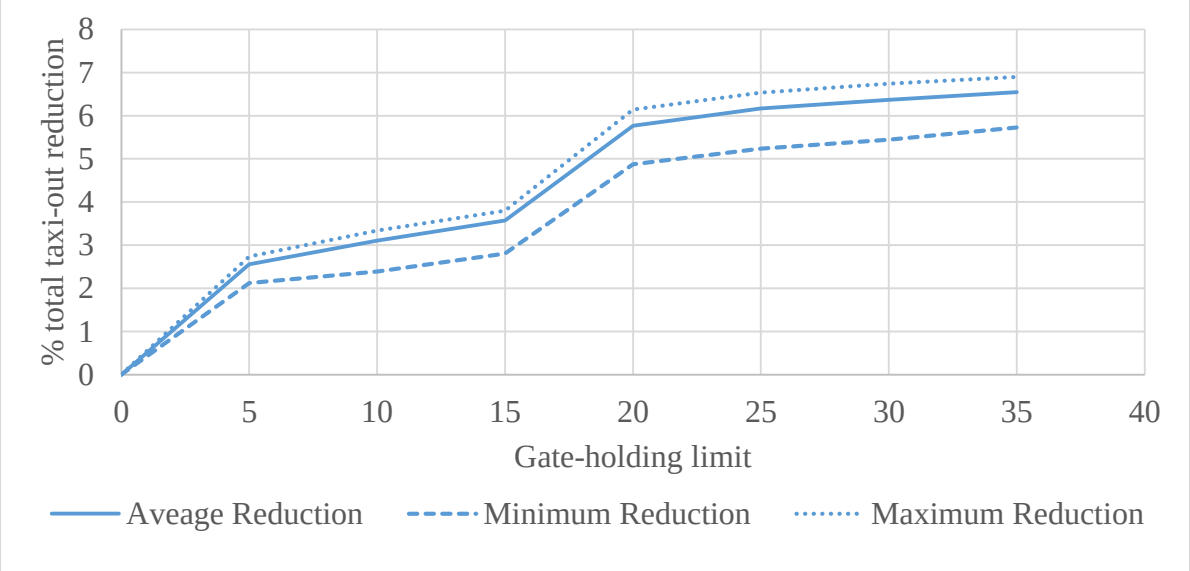

Figure 32: Variation of the percentage in taxi-out reduction with different gate-holding limits

Previous work on metering focused on the accurate development of the metering algorithms and implementation protocols. However, no previous study has proposed metering-based policies that allow operators to think of the problem from a limit-benefit exercise. Indeed, Figure 32 relates each limit to consequent benefit. Therefore, decision makers may use such information to better inform inputs and outputs of the proposed model.

One last factor should come into play while making a decision of which policy to implement: the order conservation. That is, whether the metering-based policies lead to larger violations of the First Come First Served (FCFS) order, compared to the baseline case. In order to assess such considerations, Figure 33 and Figure 34 illustrate the differences in the order in two different circumstances so as to determine the causes of order swaps.

First, Figure 33 shows order differences (or swaps) between the take-off (ATOT) order and the ready-for-pushback (EOBT) order. For all the flights in the 62-day period, we assign order number based on EOBT, we assign order based on ATOT, and then we compute the differences between these two order lists. Assigning order based on EOBT implies that such a list has not been altered by any of the metering-based policies nor by gate conflicts. However, the ordered list based on ATOT is affected by the metering policies and gate conflicts. This figure shows these differences for the eight different gate-holding strategies presented in Figure 32 so as to provide additional information during the decision.

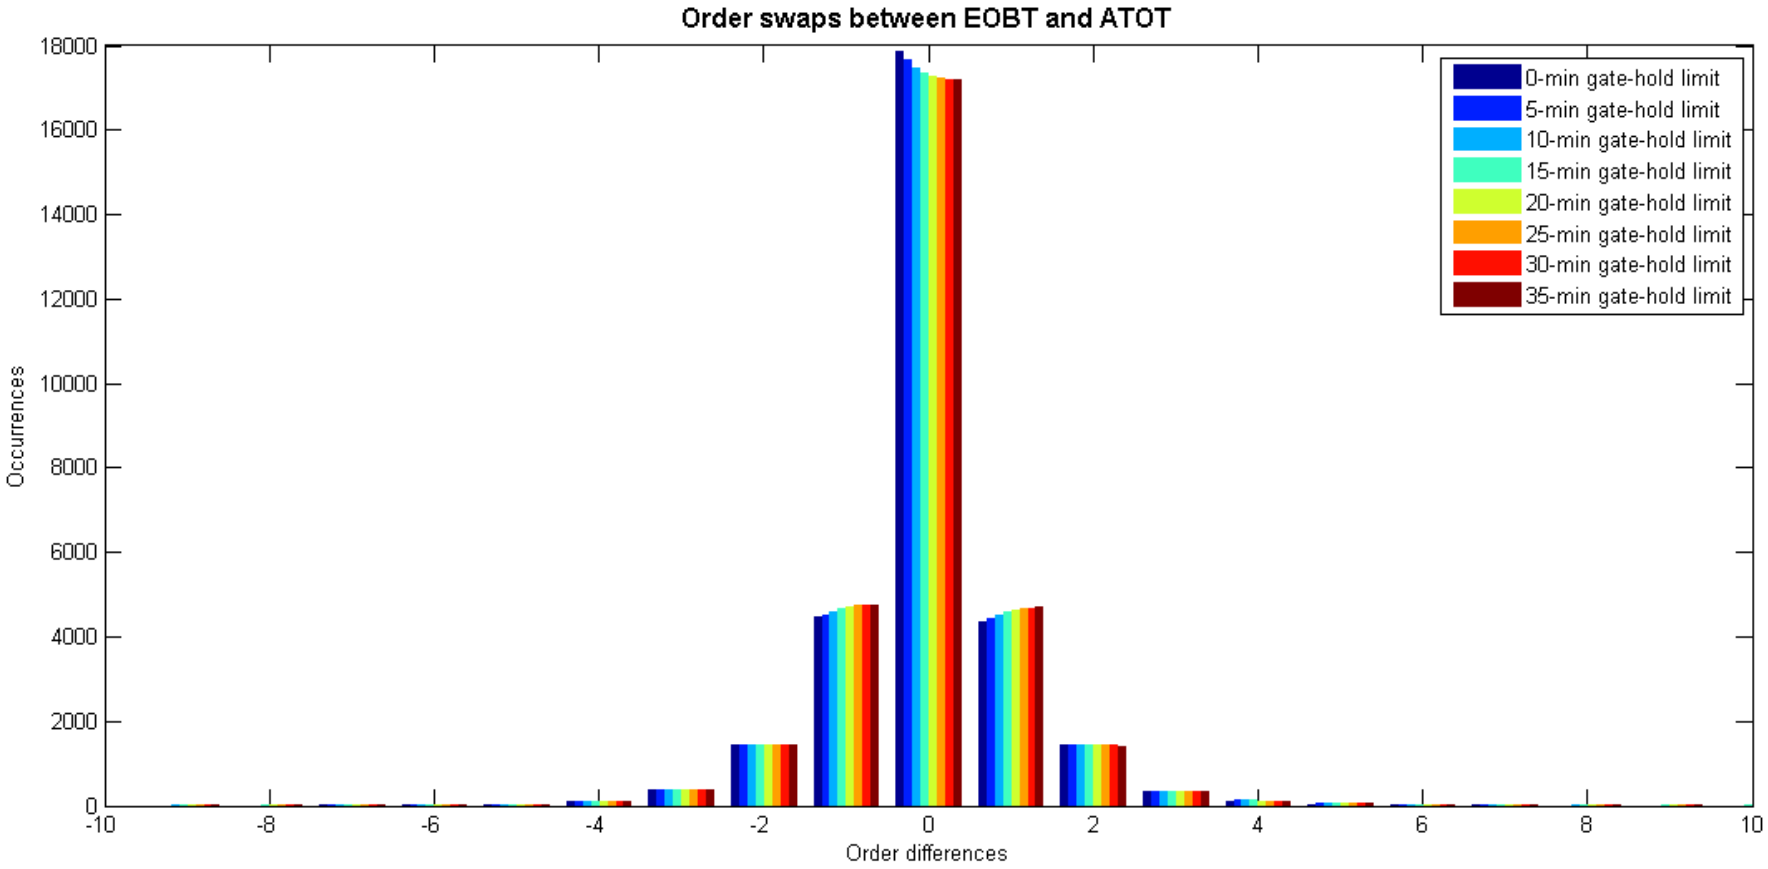

Figure 33: Order differences between the take-off order and the EOBT order (ready for push order) for the different gate-holding limit policies. Negative values correspond to flights that have been moved backward in the departure line and positive values correspond to flights that have moved forward in the departure line. This plot considers all the swaps after from the moment the aircraft is ready to push until the actual take-off time; which includes those coming from gate conflicts.

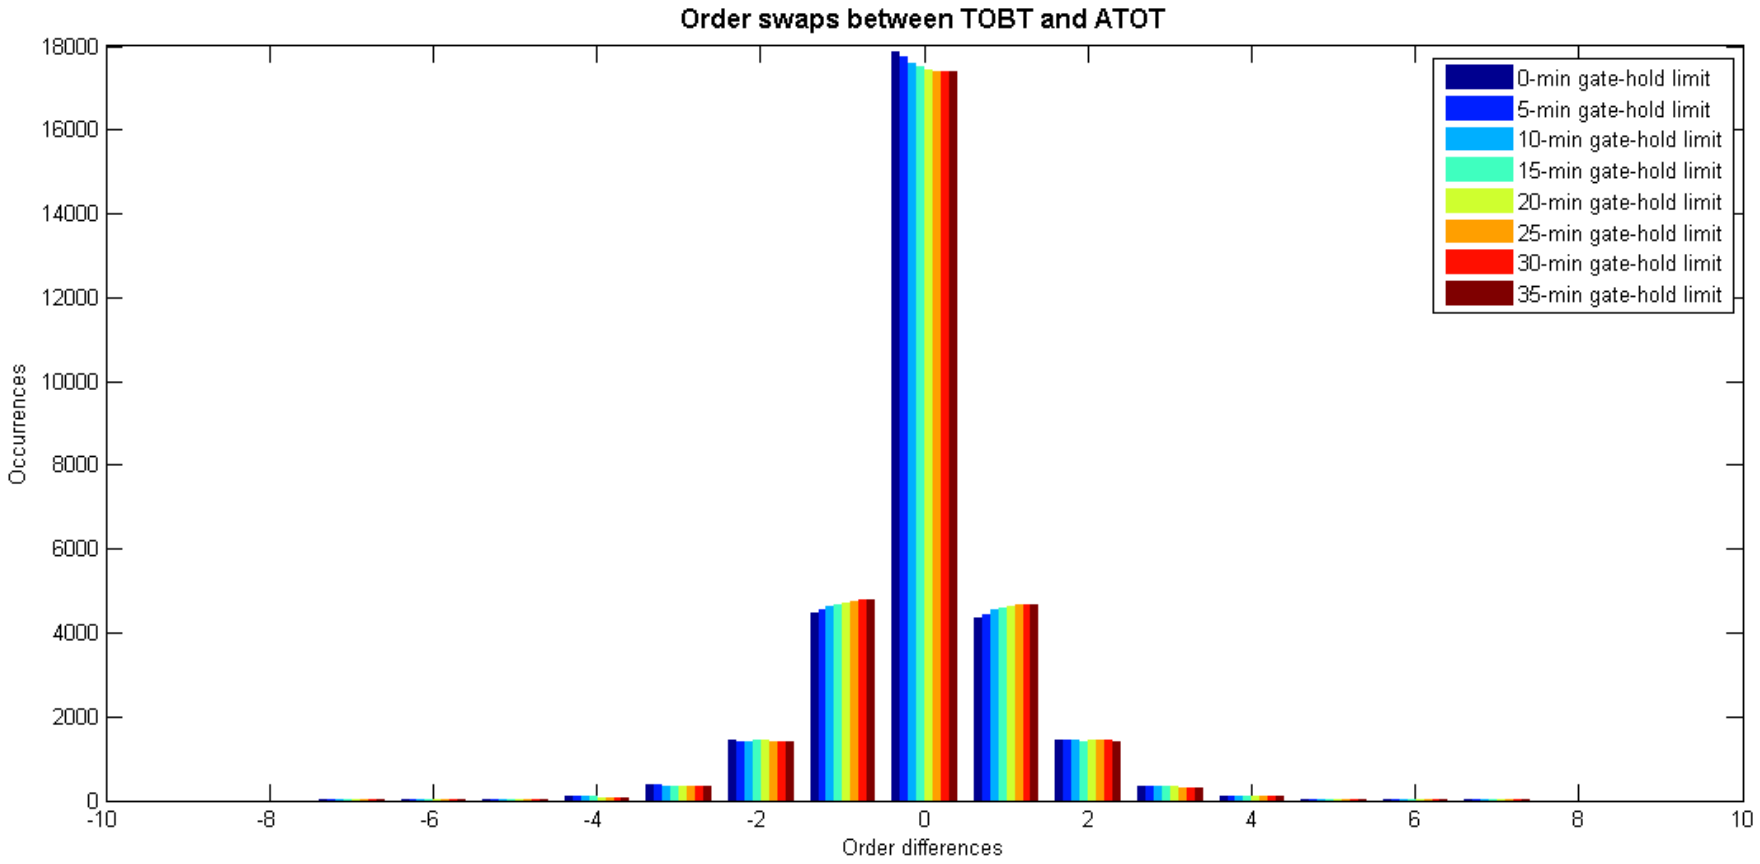

Figure 34: Order differences between the take-off order and the TOBT order (pushback order) for the different gate-holding limit policies. Negative values correspond to flights that have been moved backward in the departure line and positive values correspond to flights that have moved forward in the departure line. This plot considers all the swaps after from the moment the aircraft pushes until the actual take-off time, and thus it does not consider the swaps coming from gate conflicts.

Second, Figure 34 shows order differences between the ATOT order and the pushback (TOBT) order. For all the flights in the 62-day period, we build Figure 34 in the same way as Figure 33, but sorting for TOBT, instead of EOBT. In this case, Figure 34 also displays the differences for the eight different gate-holding policies. From the conceptual perspective, the main difference between the two figures is that in Figure 34, both the TOBT and ATOT occur after the metering-based policies and the gate conflict forced pushes have been implemented. Hence, the differences in order are due to differences in unimpeded taxi-out time. For instance, imagine the case in which aircraft 1 starts pushing back at 10:06 and has an unimpeded taxi-out time of 14 minutes, and thus reaches the departure queue at 10:20; also imagine an aircraft 2 that pushes back at 10:07, with an unimpeded taxi-out time of 12 minutes and joins the queue at 10:19, which is earlier that the time at which aircraft 1 joins the queue. In this example, aircraft 1 takes-off later than aircraft 2 even though aircraft 1 started pushing back earlier than aircraft 2. This is an example where the FCFS order is not followed. It is worth keeping the conceptual differences between Figure 33 and Figure 34 in mind to understand the former includes the effect of gate conflicts, the gate-holding limits, and the differences in unimpeded taxi-out time, whereas the latter only accounts for differences in taxi-out time. However, setting gate-holding limits applies to all flights in the same way and therefore, the order is not affected by such policies. Put another way, whereas gate-holding limits change the TOBT for some flights and reduce the chances of gate conflicts, setting limits itself does not lead to swaps because it affects all flights in the same way. Considering the rationale and explanations behind Figure 33 and Figure 34, the main observations are the following:

- A perfect policy should be able to ensure a FCFS order, which would correspond to all occurrences being represented over the 0 difference. However, such a behavior is not

observed in any of the policies. That is, all the policies fail at complying with the FCFS order.

- The darkest blue bars in both figures correspond to gate-holding limit of 0 minutes, which is equivalent to the baseline case. The presence of occurrences outside the 0 value indicate that the baseline strategy (pushback-at-discretion) is already violating the FCFS policy.
- Both figures show symmetrical distribution of differences, which align with the intuition that the number of flights moving forward should balance out with those moving backward.
- The policies with a larger gate-holding limit lose some of occurrences compliant with the FCFS policy toward +1 or -1. That is, the three central bins - corresponding to -1, 0, and +1- are the only bins experiencing noticeable changes, in both figures. The interesting part is that what is being lost in the 0 bin is gained in both the -1 and +1, which makes sense because if an aircraft moves back, then another one moves forward. Therefore, losing one 0-bin occurrence leads to one +1 observation and a -1 observation.
- The decreasing trend for the 0-bin and the increasing trend for bins -1 and +1 are slightly more pronounced in Figure 33 than in Figure 34. Recalling that differences in Figure 33 account for gate conflicts and unimpeded taxi-out time, and Figure 34 accounts for unimpeded taxi-out time, we can state that gate conflicts are leading to slightly more pronounced change in order trend. However, comparing both figures, it is also clear that the biggest part of the increasing/decreasing trend is caused by the differences in unimpeded taxi-out time, which are triggered by the changed in TOBT. Put another way, metering based policies lead to new TOBT that can easily translate into different surface dynamics. In particular, the different terminal-dependent unimpeded taxi-out times are the responsible for differences in take-off order. Notwithstanding these differences in TOBTs

and unimpeded taxi-out time, the impact on the simulation is relatively small, as these switches are of 1 position forward and backward.

After understanding that metering policies lead to results that are not significantly different from those in the baseline at an aggregate level, the next step is to assess whether there are differences at an airline level. To this end, Figure 33 plots the order differences between EOBT and ATOT. For comparison clarity purposes, Figure 35, on the one hand, only plots the baseline case and the unrestricted case given that they are the most extreme ones and provide upper and lower bounds values; on the other hand, Figure 35 plots percentage values of occurrences, as opposed to absolute number of occurrences, in order to facilitate the comparison exercise (different airlines operate different numbers of departures and thus, it is not fair to compare the absolute number of occurrences).

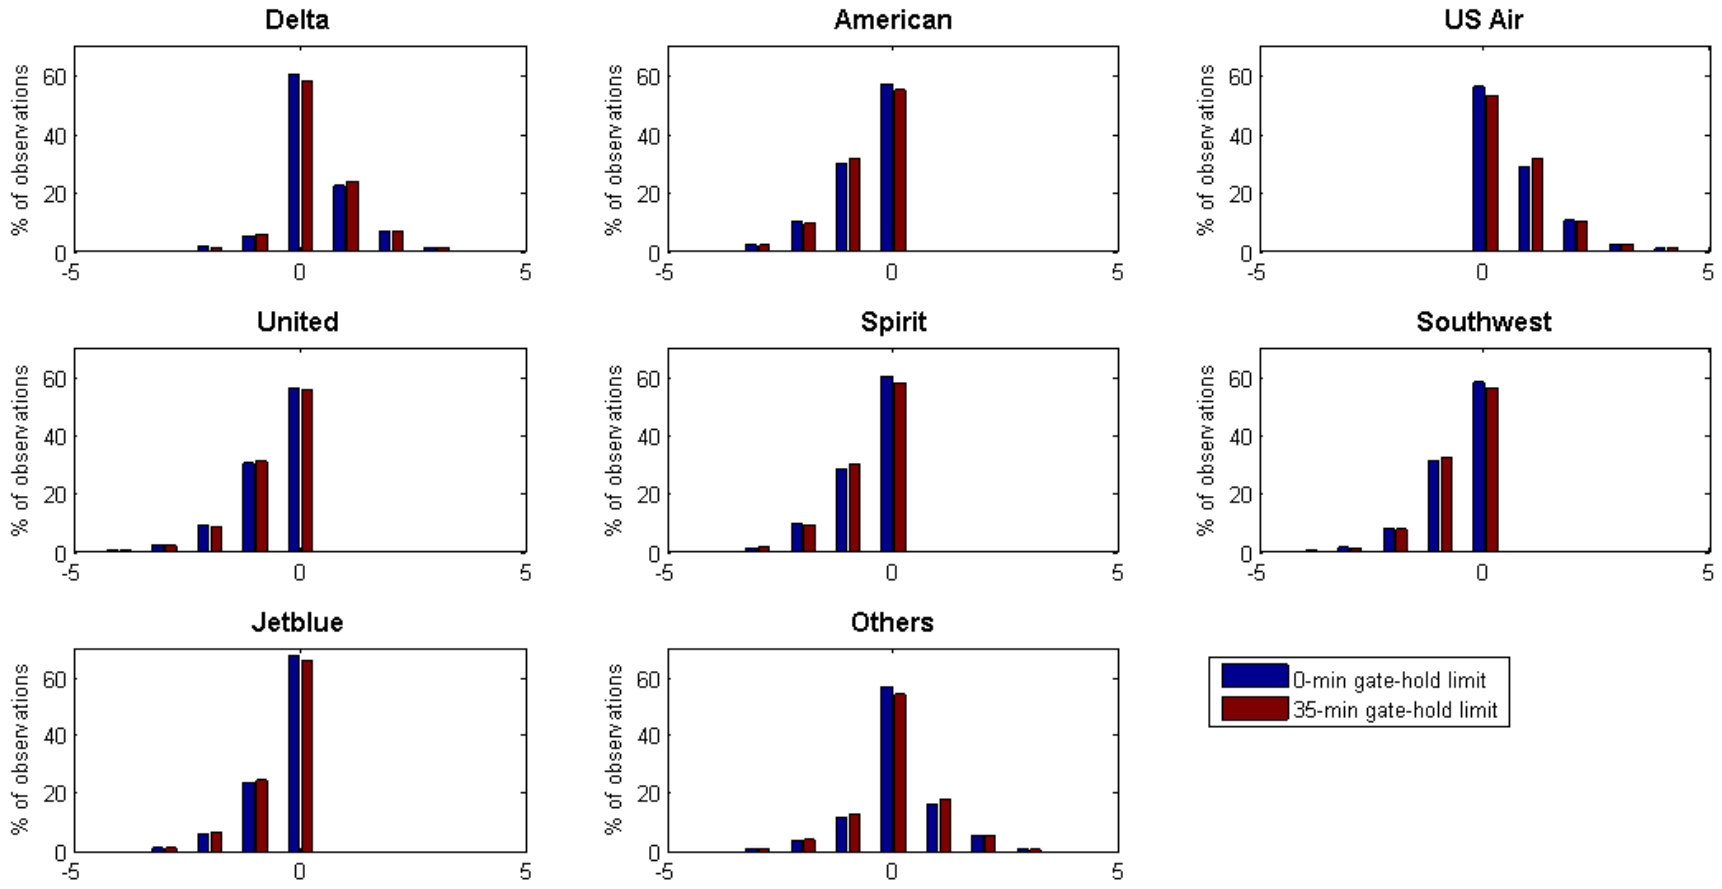

Figure 35: Order differences between the take-off order and the EOBT order for the different airlines and for the two extreme gate-holding policies. Negative values correspond to flights that have been moved backward in the departure line and positive values correspond to flights that have moved forward in the departure line.

The results in Figure 35 support our hypothesis that most of the order differences are caused by the difference in taxi-out time by terminal. Delta and US Airways (operating in Terminals C and D) present a similar behavior; indeed, these carriers tend to see their flights either complying with the FCFS (equivalent to difference equal to 0) policy or moved forward. Conversely, the rest of the carriers (American, United, Spirit, Southwest, and JetBlue), operating in Terminal B, see their flights either complying with the FCFS order or moved backward. Although these findings are not ideal, the most relevant take-away from Figure 35 is that metering barely worsens the situation in this regard, given that variations in the bins above are marginal. If the differences between the most extreme gate-holding policies are small, any gate-holding policy in between will certainly present a portion of it, which will be, indeed, small.

This chapter has proven that metering is effective at reducing taxi-out times without causing a significant increase in gate conflicts, while barely altering the departure order, avoiding taxi-out delays due to gate-holding, and ensuring fair distribution of benefits (proportional to departure share). Finally, delving deeper into the dynamics of gate-holding policies, the percentage reduction in taxi-out time evolves in a non-linear fashion with the gate-holding limits, which provides stakeholders with a portfolio of implementation policies connecting operational effort to monetary benefits. Such a finding can help ease discussions on the implementation of metering-based strategies aiming at reducing airport surface congestion, which offer emission and fuel burn saving opportunities.

4. CONCLUSIONS

4.1. SUMMARY AND POLICY IMPLICATIONS

This thesis presents modifications to the metering strategy consisting of limiting the maximum gate-holding time. The main goal of this thesis was to assist decision makers in the process of implementing metering and to ease the practical implementation challenges that arise at a chronically congested airport such as LaGuardia. To this end, this thesis introduces a portfolio of metering-based policies, and a tool to evaluate benefits and costs of implementing one of the policies. On the benefits side, the thesis looks at taxi-out time reductions up to 6.7%, which lead to reductions in fuel burn and emissions. On the challenges side, the thesis analyses gate-holding time, gate conflicts and order differences; while the gate-holding time becomes significant with metering, gate conflicts and order swaps increase only slightly.

As for the benefits of implementing metering-based policies, they follow a non-linear increasing function of the gate-holding limit. The increasing portion of the function supports that the larger the gate-holding limit, the larger the taxi-out reduction limits. However, the non-linear portion of the function requires decision makers to use the tools and analysis proposed in this thesis to pick the best policy.

From a fairness perspective, all the policies lead to gate-holding times that are commensurate with the reduction in taxi-out times, which indicates that every holding minute translates into a minute of taxi-out reduction. In addition, the gate-holding times and taxi-out time saving in all policies are allocated in proportion to the number of departure operations amongst the different air carriers. Taking all these results into account, airlines can do an internal analysis of costs associated with implementing each strategy and compare these costs with the benefits. Then, based on individual

airline cost-benefit trade-offs, all the stake-holders should set an implementation plan, deciding which, if any, stages would help overcome some of the challenges, enabled by the policy portfolio's flexibility.

4.2. FUTURE RESEARCH

The final part of the thesis suggests directions for further investigation. As presented in Section 2.3.2, the major trade-offs while deciding the time period length of the simulation are computational cost, prediction accuracy, usefulness during implementation and synchronization with the model structure. However, prediction accuracy was the variable that led us to pick 15 minutes as time period, because of the suboptimal results of the current model for time periods of 30 and 60 minutes. Figure 33, Figure 34, and Figure 35 depict the prediction errors for time period durations of 15, 30, and 60 minutes.

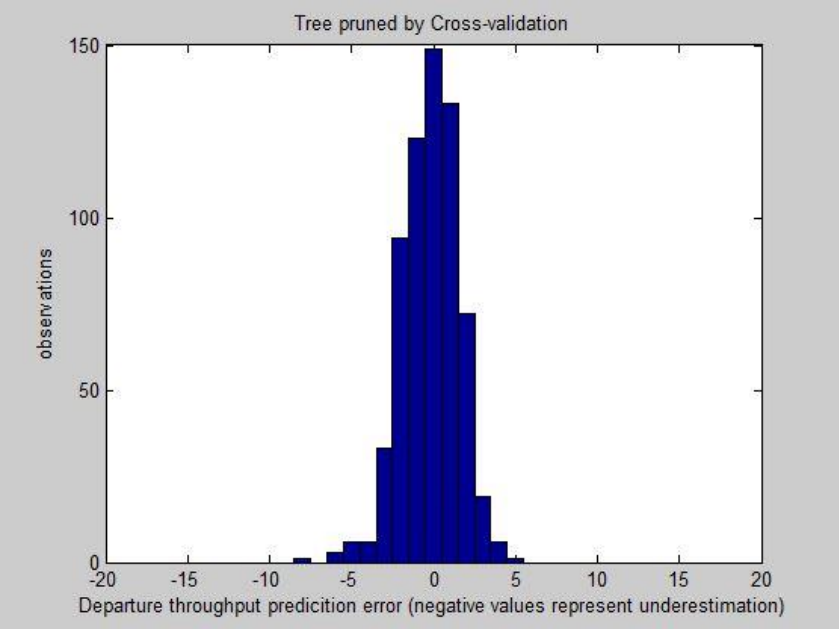

Figure 36: Prediction error with the estimated regression trees with a model time period of 15 minutes.

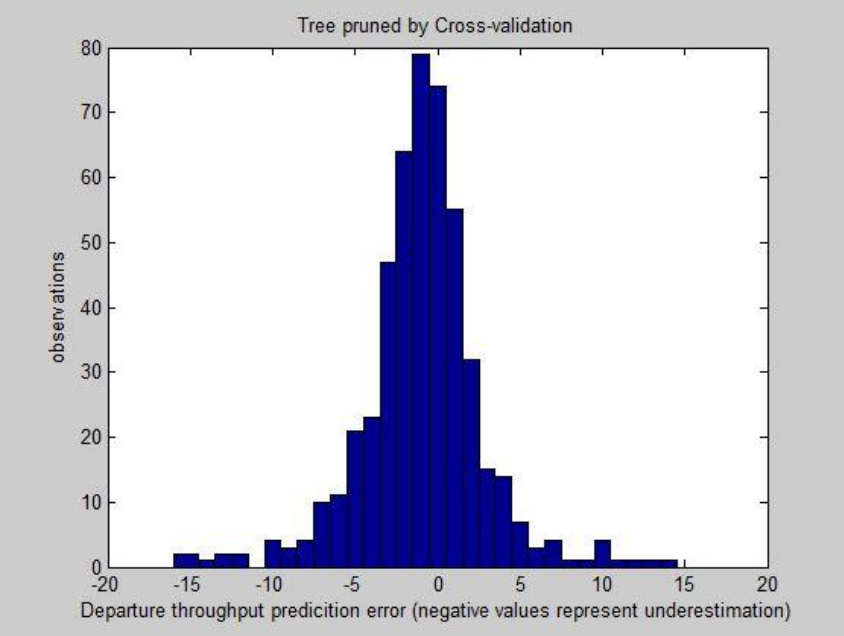

Figure 37: Prediction error with the estimated regression trees with a model time period of 30 minutes.

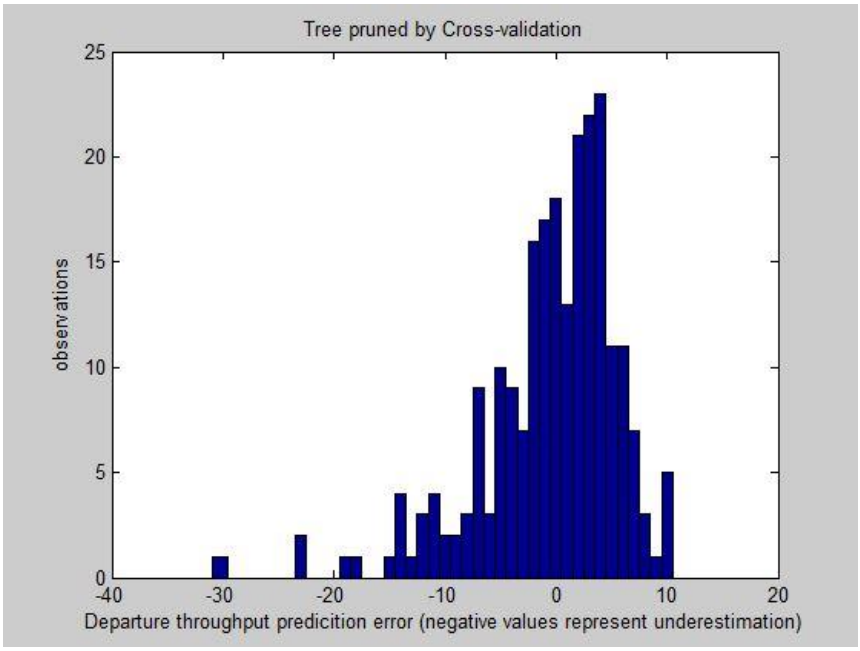

Figure 38: Prediction error with the estimated regression trees with a model time period of 60 minutes.

Based on these observations, future research could study more sophisticated and dynamic models that can improve predictions for periods longer than 15 minutes. Previous studies (Simaiakis 2013) suggest using Dynamic Programming (DP) algorithms to better track surface traffic and departing queuing performance. DP can be implemented in a similar way to what is described in this thesis for metering-based strategies. DP algorithms also hold aircraft at the gate to achieve reductions in taxi-out time. Should DP algorithms be successful at reducing prediction error, the time period length could be used as another factor during the implementation design process, in addition to the gate-holding limit.

I. APPENDIX A: UNIMPEDED TAXI-OUT TIME PLOTS

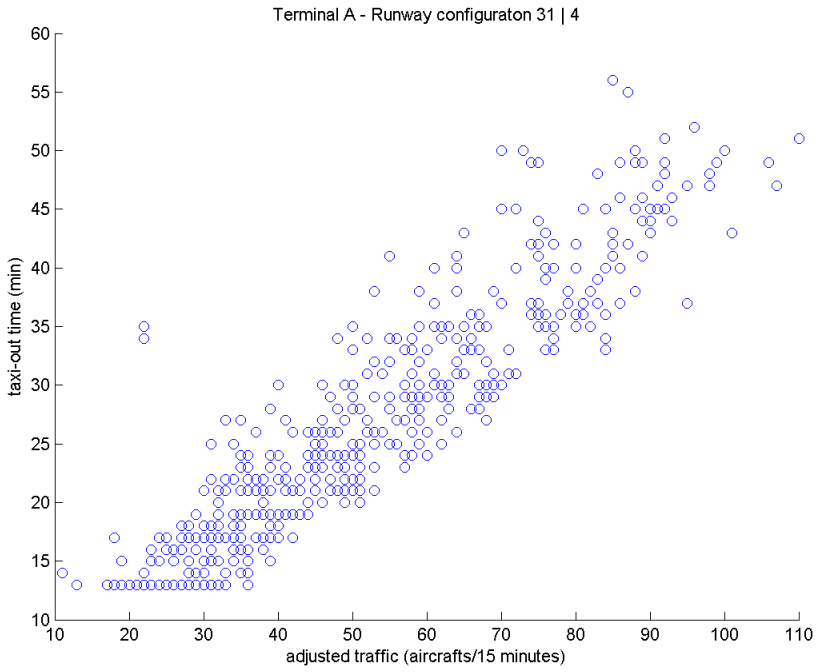

Figure A. I: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal A with runway configuration 31|4. Data from July 1, 2013 to August 30, 2013

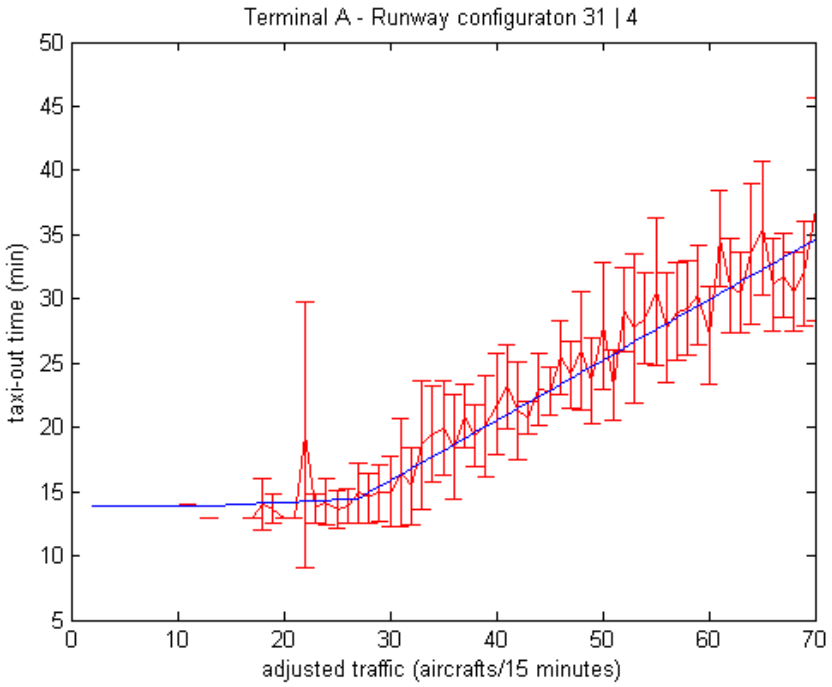

Figure A. II: Mean, standard error, and fit function for flights from Terminal A when the airport operates with runway configuration 31|4. Data from July 1, 2013 to August 30, 2013

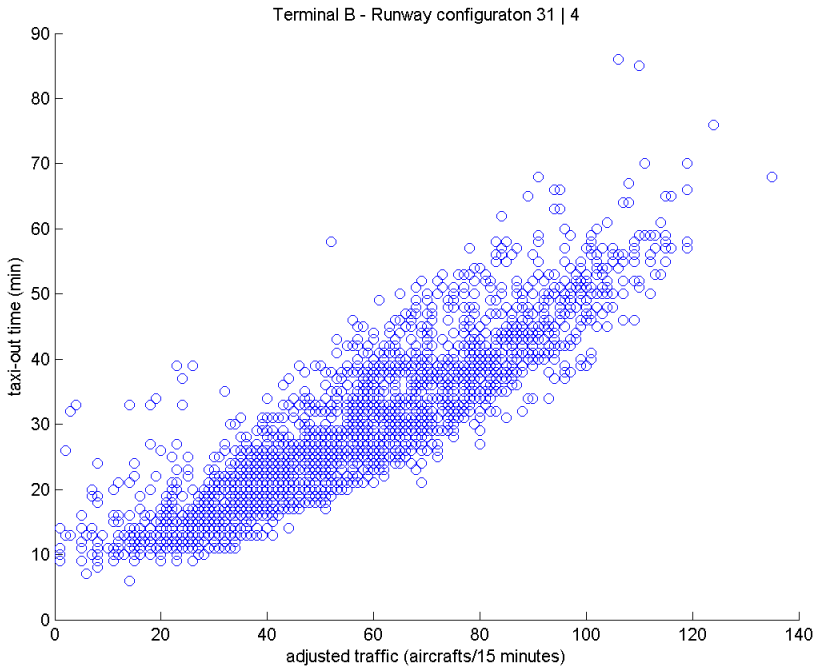

Figure A. III: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal B with runway configuration 31|4. Data from July 1, 2013 to August 30, 2013

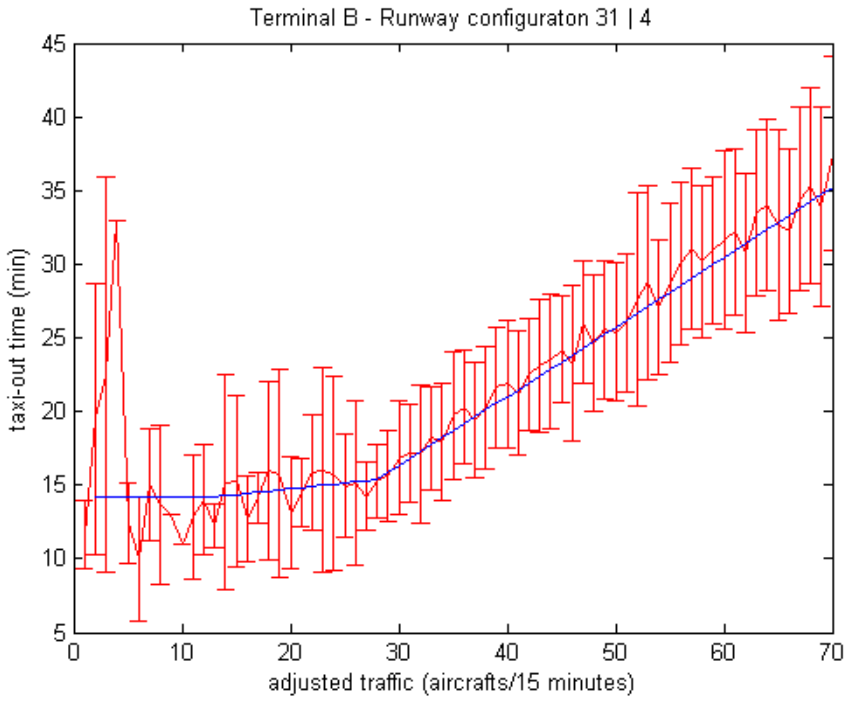

Figure A. IV: Mean, standard error, and fit function for flights from Terminal B when the airport operates with runway configuration 31|4. Data from July 1, 2013 to August 30, 2013

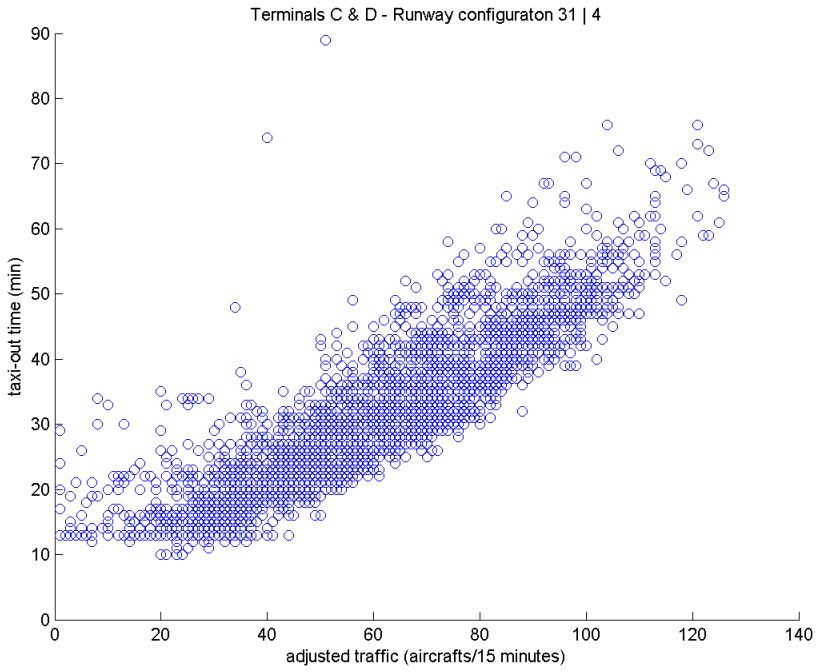

Figure A. V: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from Terminals C&D with runway configuration 31|4. Data from July 1, 2013 to August 30, 2013

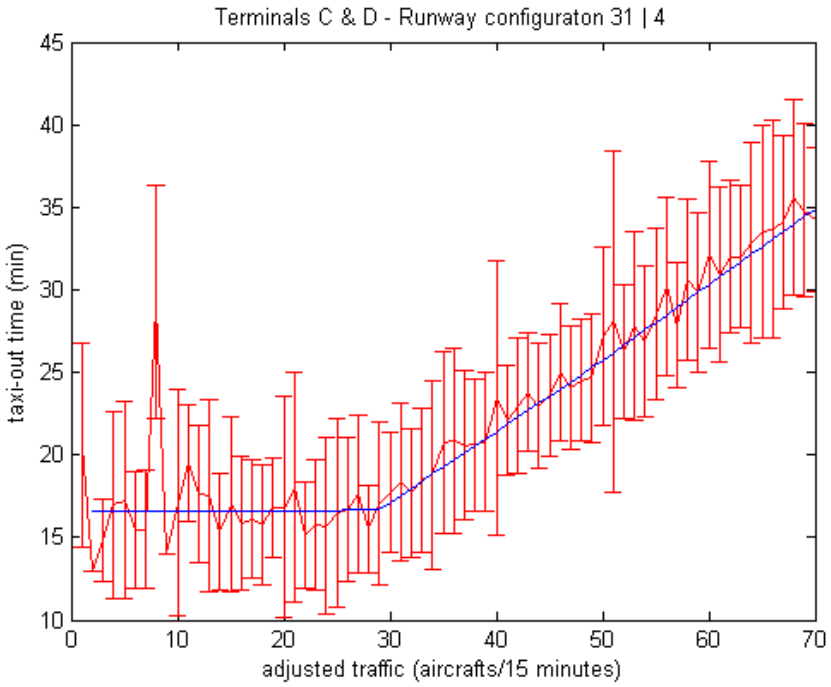

Figure A. VI: Mean, standard error, and fit function for flights from Terminals C&D when the airport operates with runway configuration 31|4. Data from July 1, 2013 to August 30, 2013

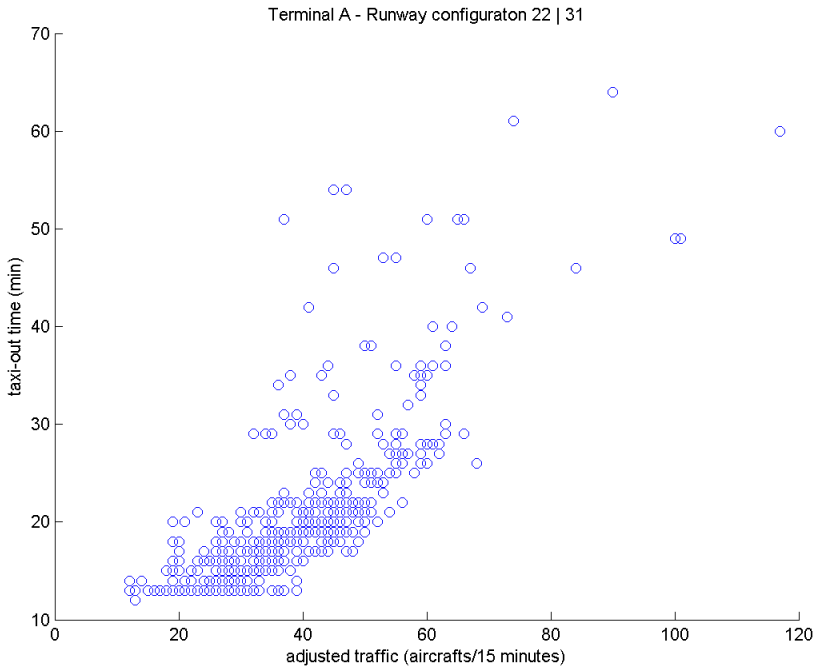

Figure A. VII: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal A with runway configuration 22|31. Data from July 1, 2013 to August 30, 2013

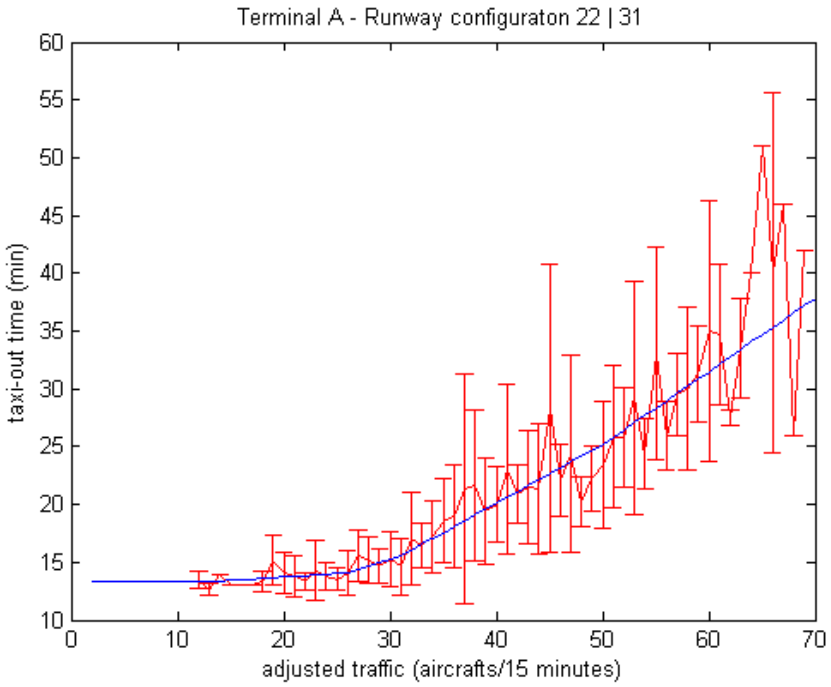

Figure A. VIII: Mean, standard error, and fit function for flights from Terminal A when the airport operates with runway configuration 22|31. Data from July 1, 2013 to August 30, 2013

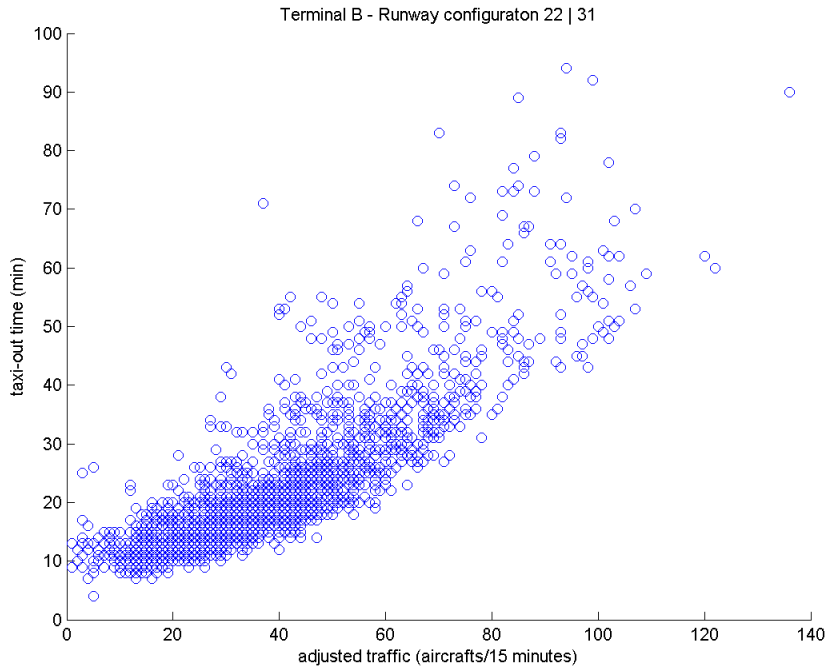

Figure A. IX: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal B with runway configuration 22|31. Data from July 1, 2013 to August 30, 2013

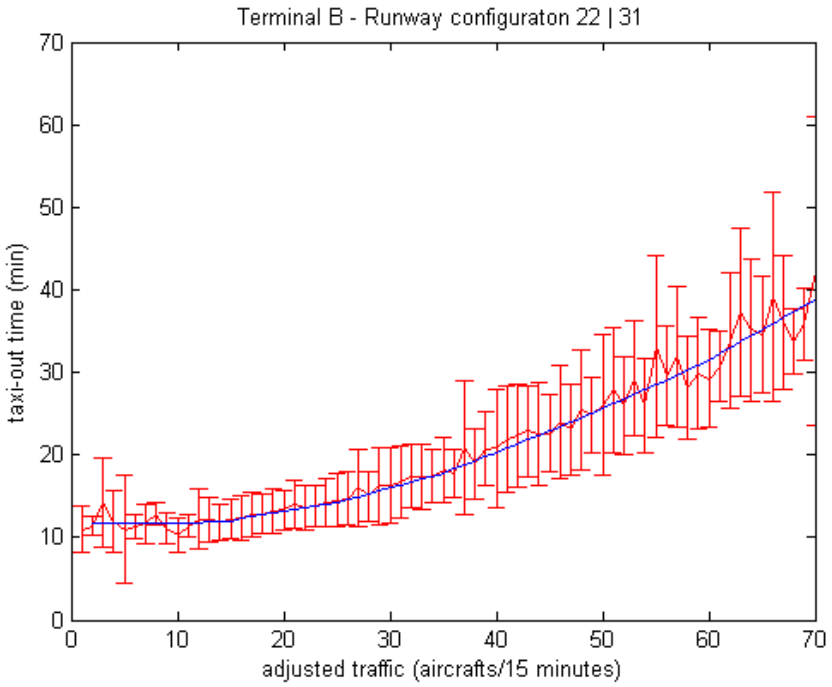

Figure A. X: Mean, standard error, and fit function for flights from Terminal B when the airport operates with runway configuration 22|31. Data from July 1, 2013 to August 30, 2013

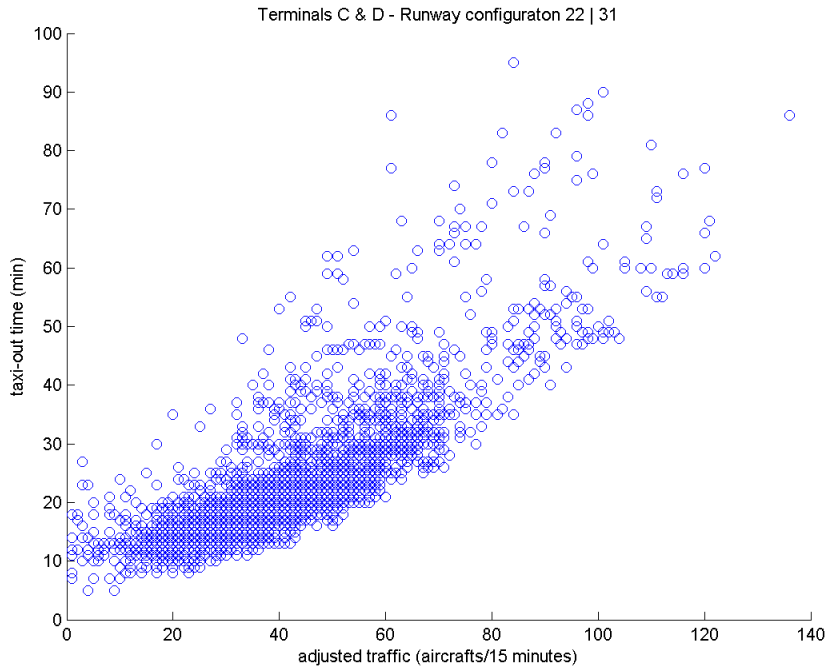

Figure A. XI: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from Terminals C&D with runway configuration 22|31. Data from July 1, 2013 to August 30, 2013

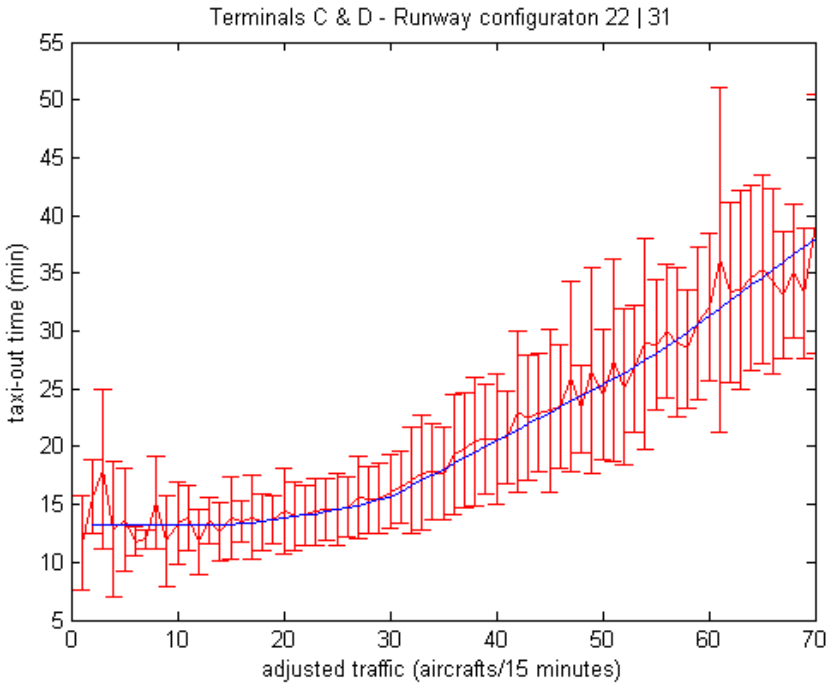

Figure A. XII: Mean, standard error, and fit function for flights from Terminals C&D when the airport operates with runway configuration 22|31. Data from July 1, 2013 to August 30, 2013

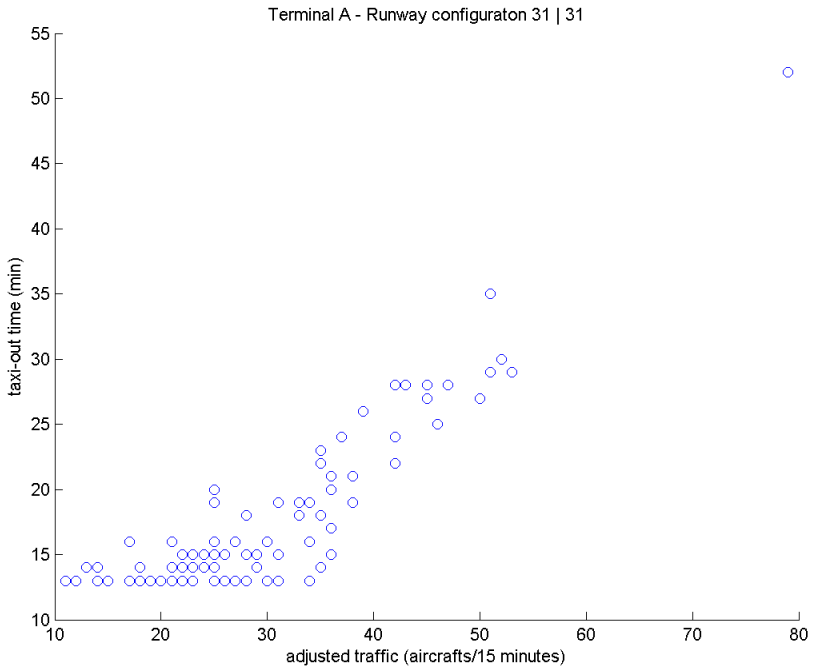

Figure A. XIII: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal A with runway configuration 31|31. Data from July 1, 2013 to August 30, 2013

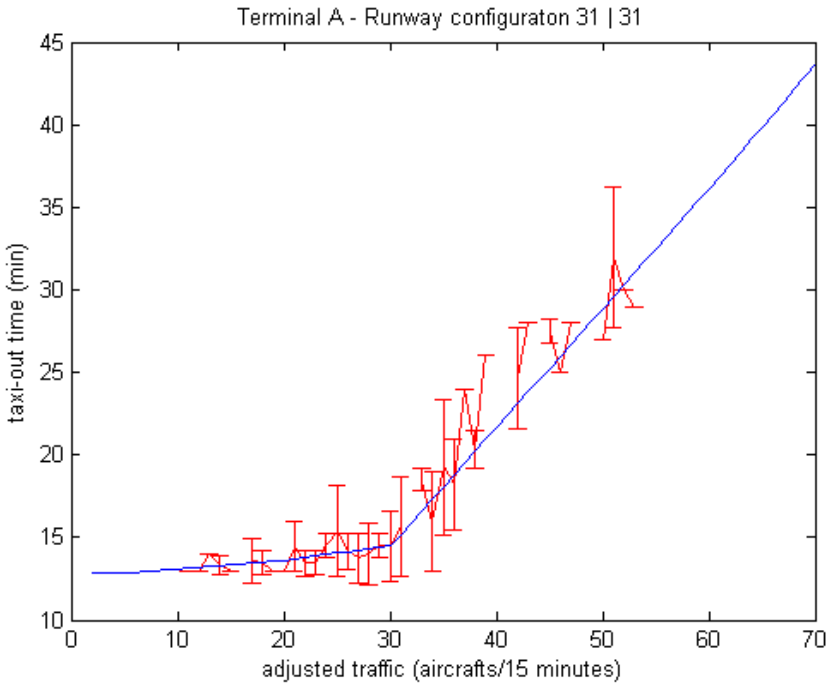

Figure A. XIV: Mean, standard error, and fit function for flights from Terminal A when the airport operates with runway configuration 31|31. Data from July 1, 2013 to August 30, 2013

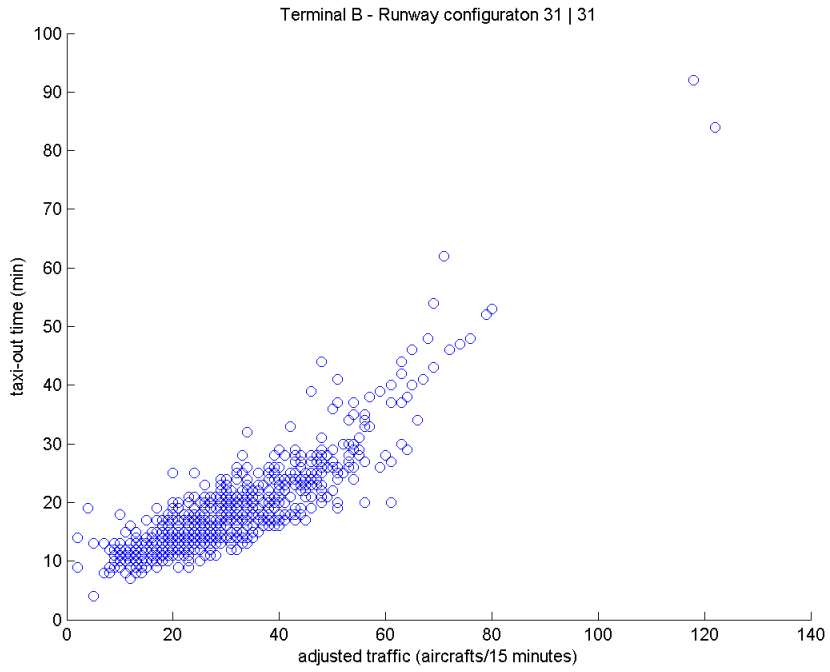

Figure A. XV: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal B with runway configuration 31|31. Data from July 1, 2013 to August 30, 2013

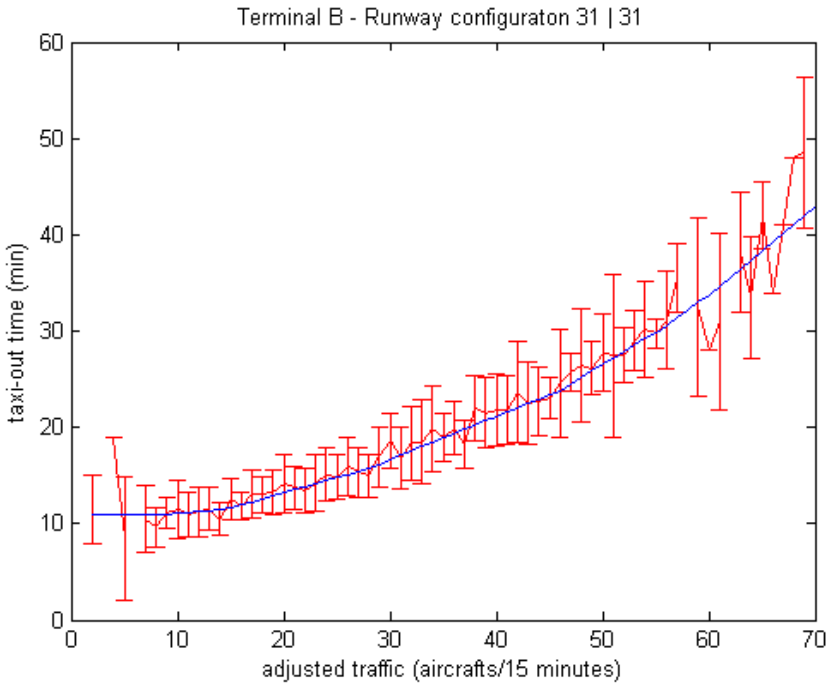

Figure A. XVI: Mean, standard error, and fit function for flights from Terminal B when the airport operates with runway configuration 31|31. Data from July 1, 2013 to August 30, 2013

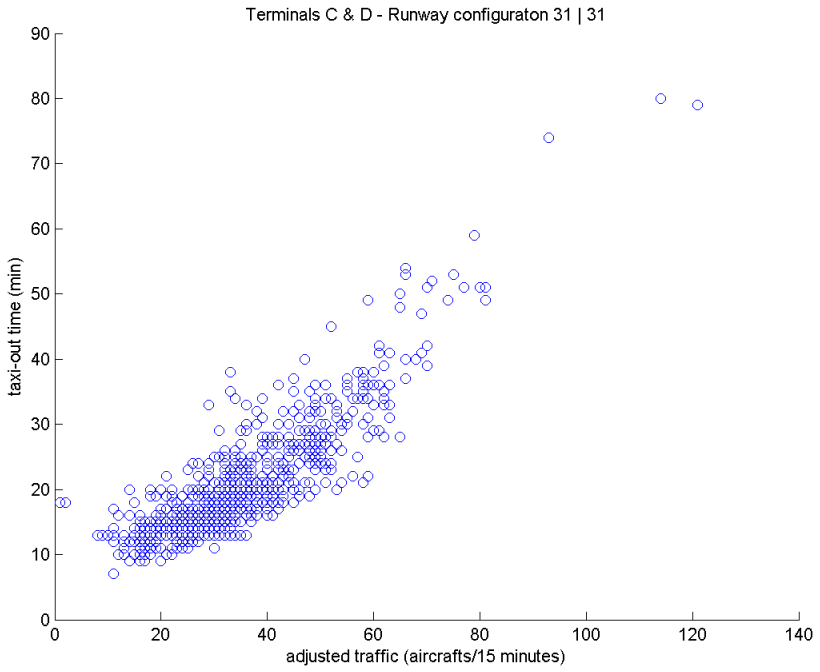

Figure A. XVII: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from Terminals C&D with runway configuration 31|31. Data from July 1, 2013 to August 30, 2013

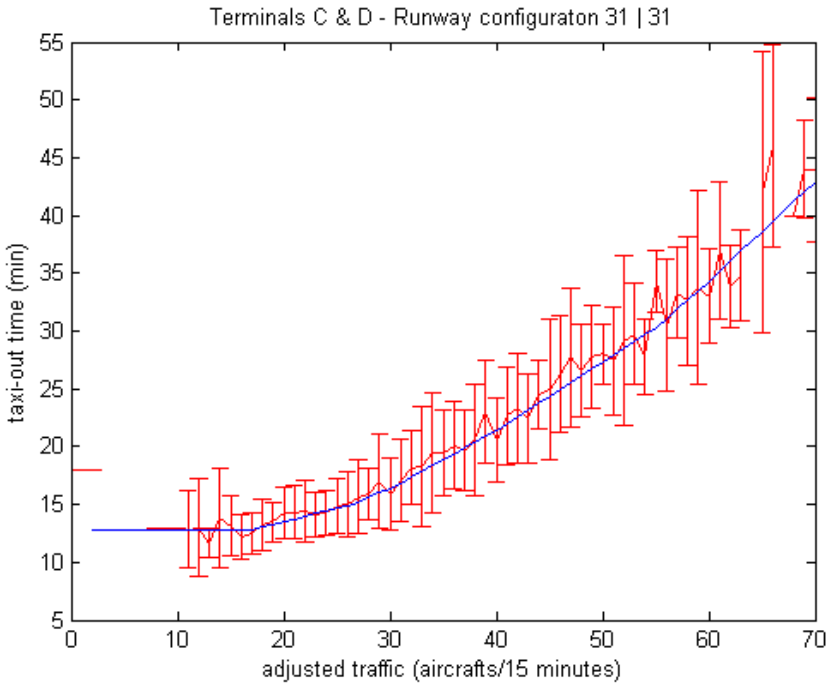

Figure A. XVIII: Mean, standard error, and fit function for flights from Terminals C&D when the airport operates with runway configuration 31|31. Data from July 1, 2013 to August 30, 2013

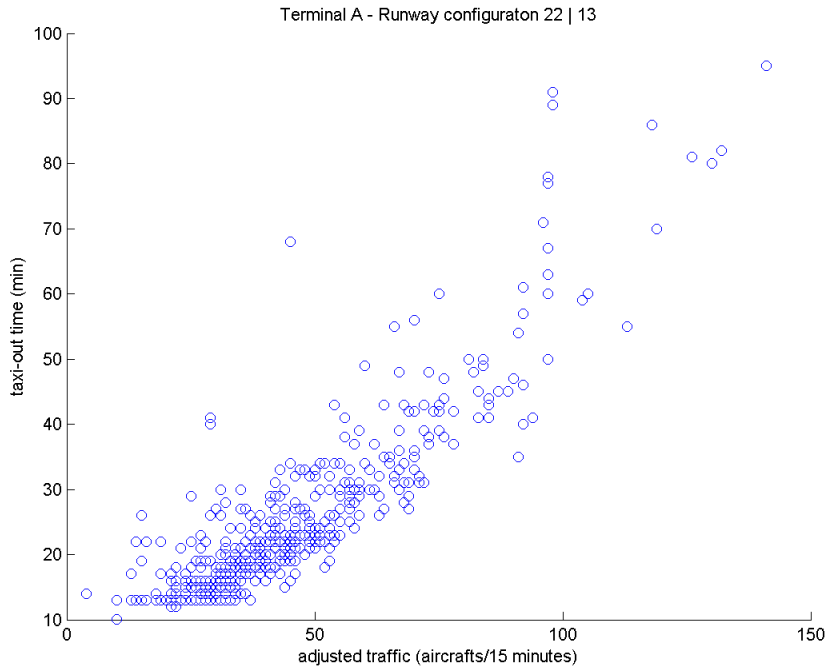

Figure A. XIX: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal A with runway configuration 22|13. Data from July 1, 2013 to August 30, 2013

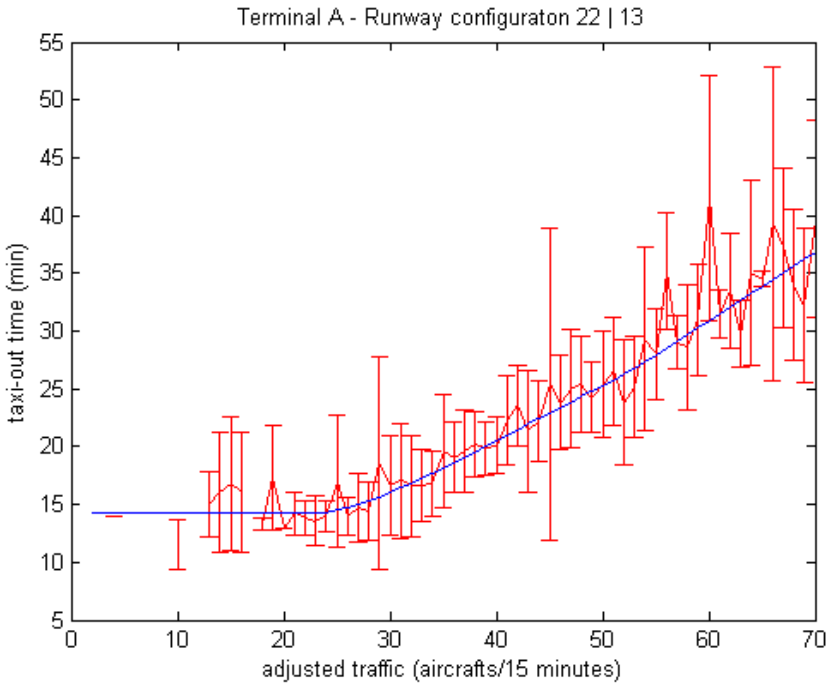

Figure A. XX: Mean, standard error, and fit function for flights from Terminal A when the airport operates with runway configuration 22|13. Data from July 1, 2013 to August 30, 2013

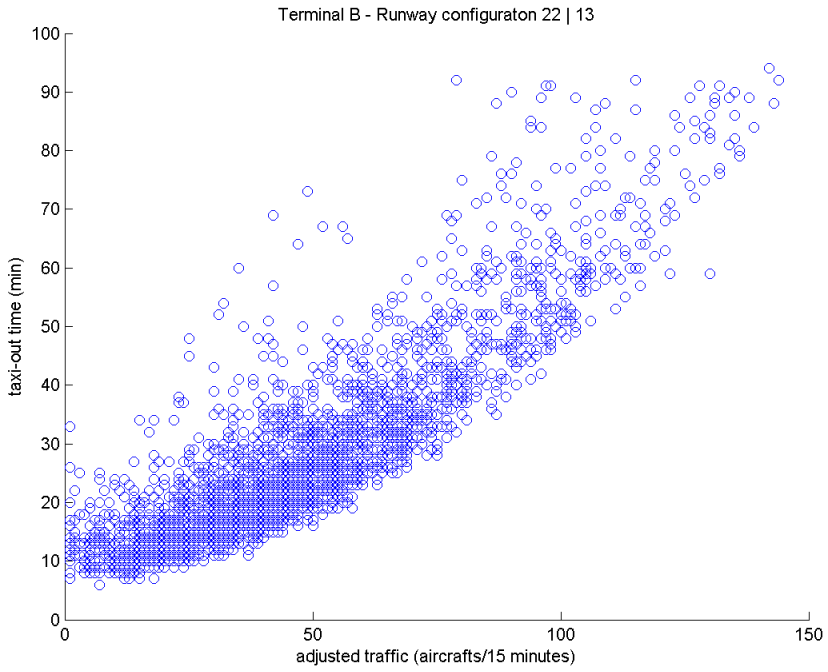

Figure A. XXI: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal B with runway configuration 22|13. Data from July 1, 2013 to August 30, 2013

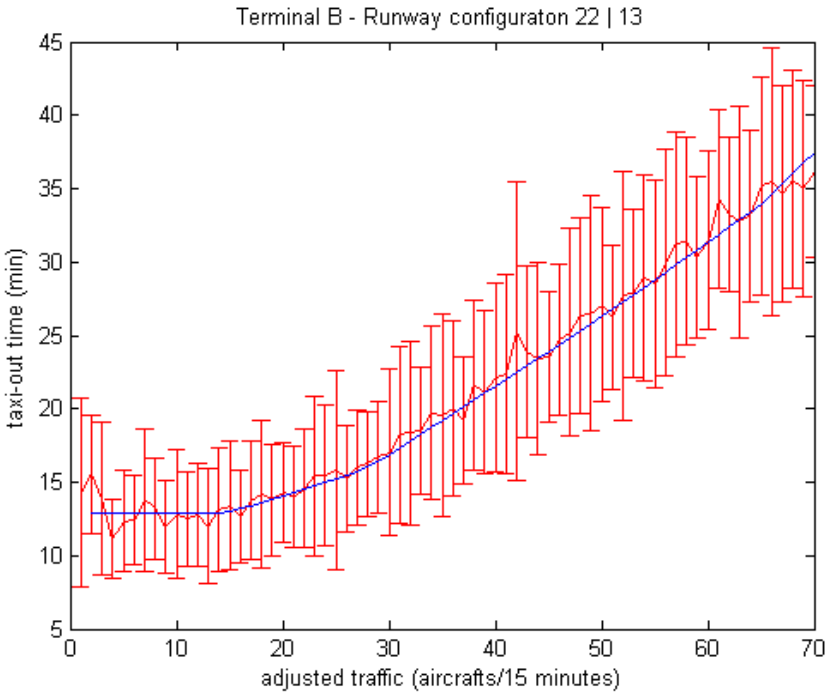

Figure A. XXII: Mean, standard error, and fit function for flights from Terminal B when the airport operates with runway configuration 22|13. Data from July 1, 2013 to August 30, 2013

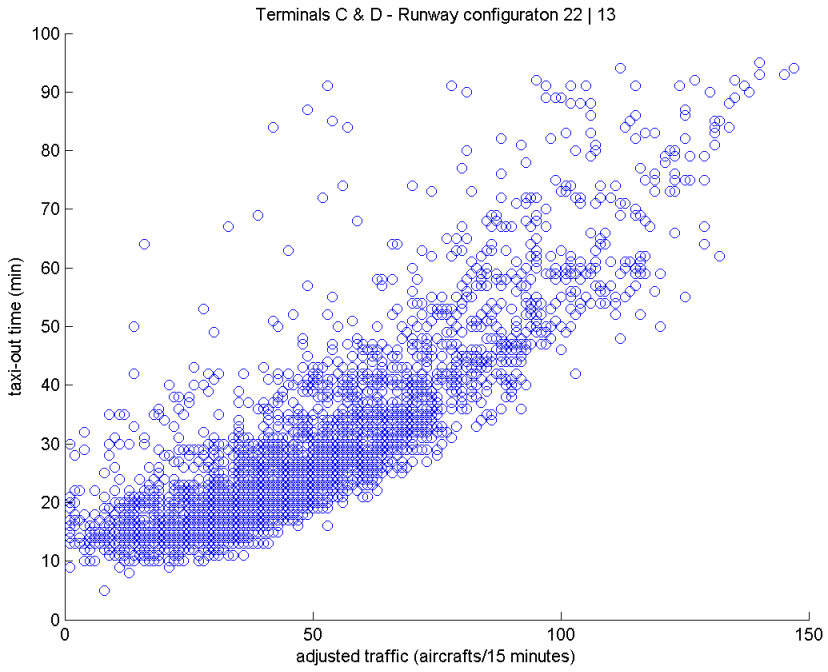

Figure A. XXIII: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from Terminals C&D with runway configuration 22|13. Data from July 1, 2013 to August 30, 2013

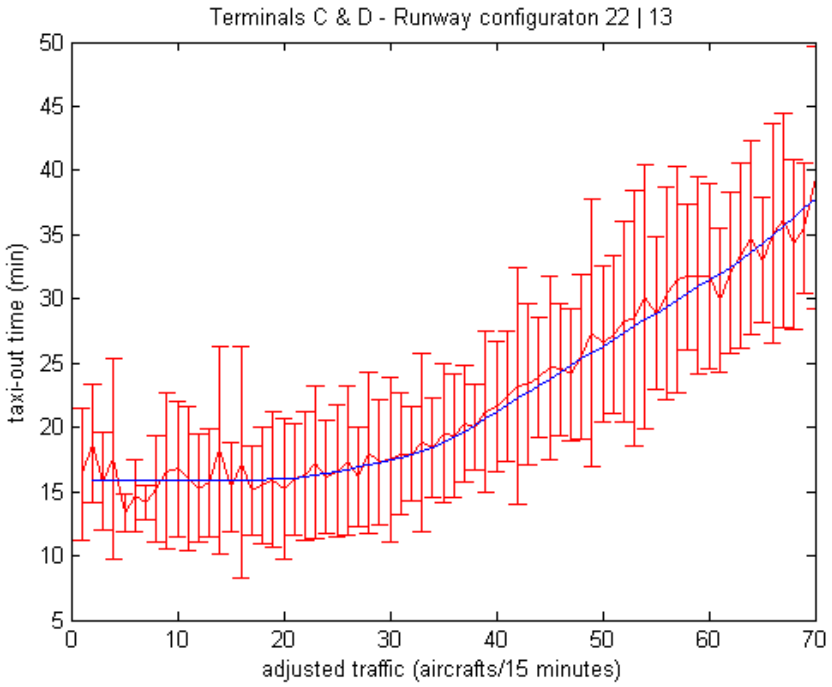

Figure A. XXIV: Mean, standard error, and fit function for flights from Terminals C&D when the airport operates with runway configuration 22|13. Data from July 1, 2013 to August 30, 2013

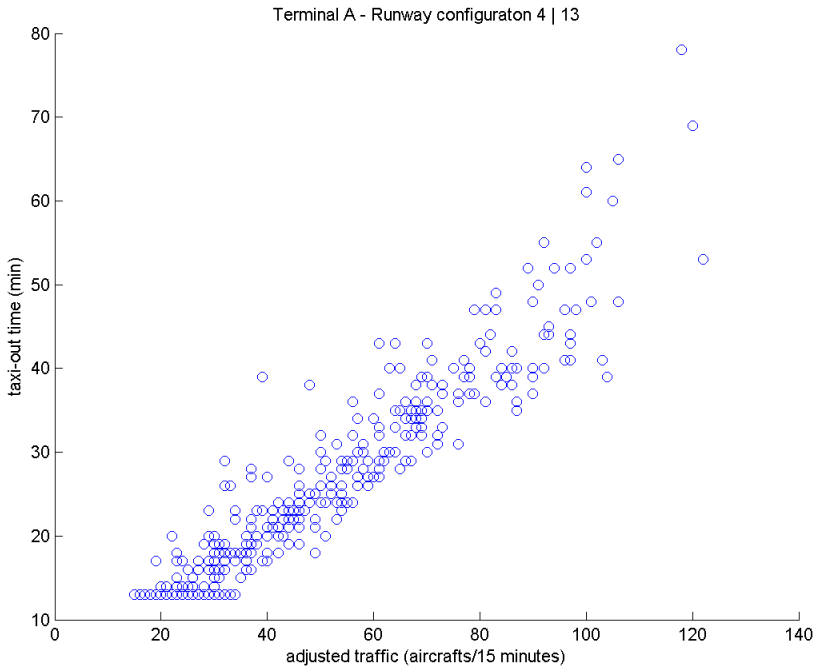

Figure A. XXV: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal A with runway configuration 4|13. Data from July 1, 2013 to August 30, 2013

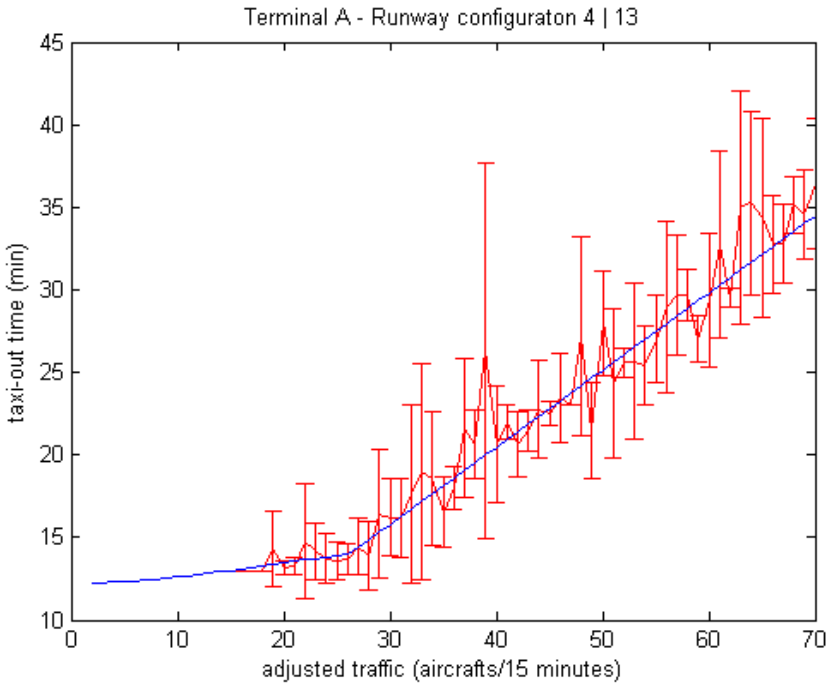

Figure A. XXVI: Mean, standard error, and fit function for flights from Terminal A when the airport operates with runway configuration 4|13. Data from July 1, 2013 to August 30, 2013

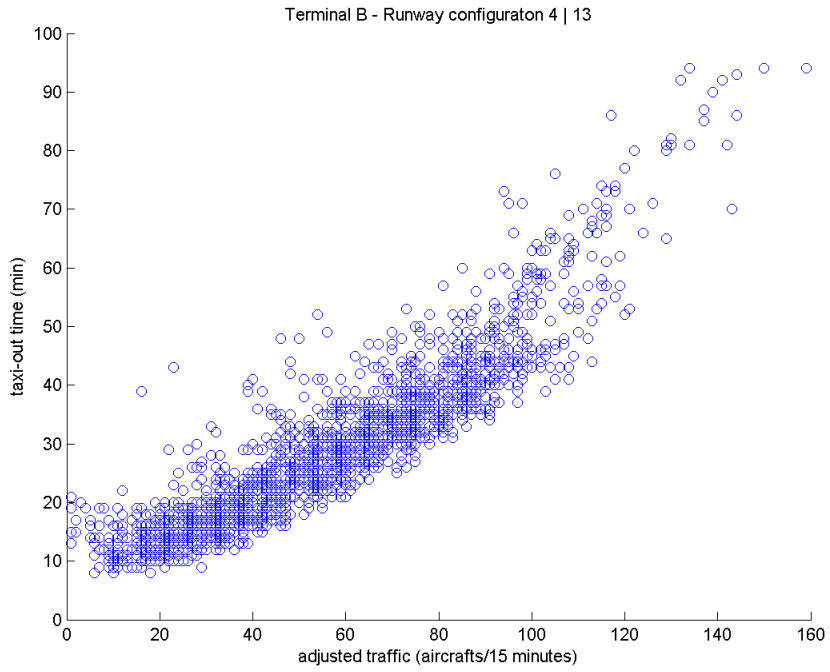

Figure A. XXVII: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from terminal B with runway configuration 4|13. Data from July 1, 2013 to August 30, 2013

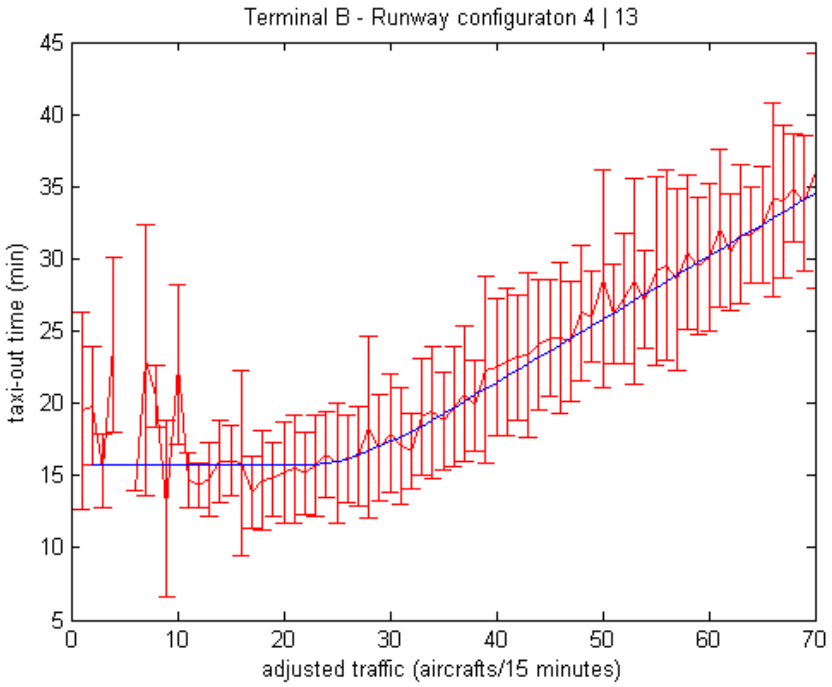

Figure A. XXVIII: Mean, standard error, and fit function for flights from Terminal B when the airport operates with runway configuration 4|13. Data from July 1, 2013 to August 30, 2013

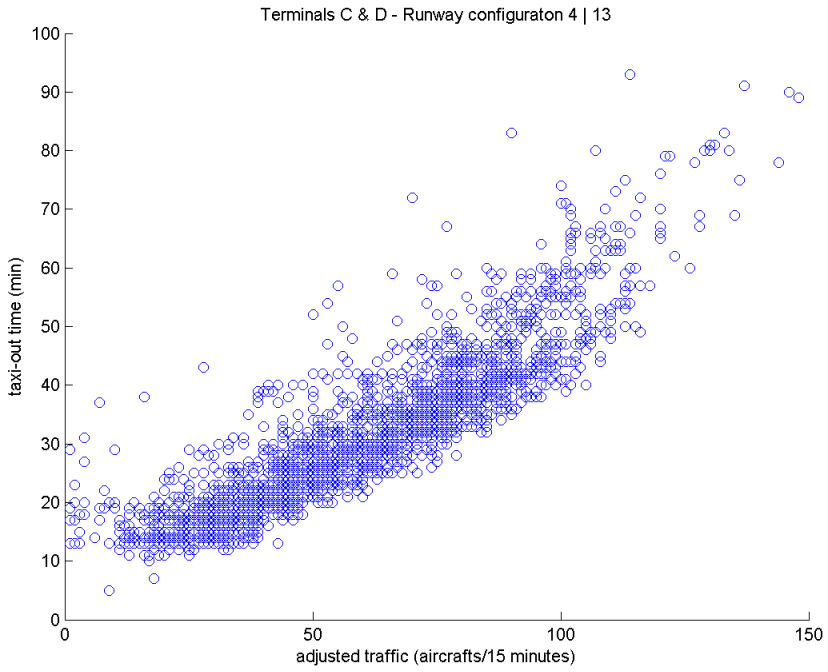

Figure A. XXIX: Scatter plot with taxi-out time as a function of the adjusted traffic for flights from Terminals C&D with runway configuration 4|13. Data from July 1, 2013 to August 30, 2013

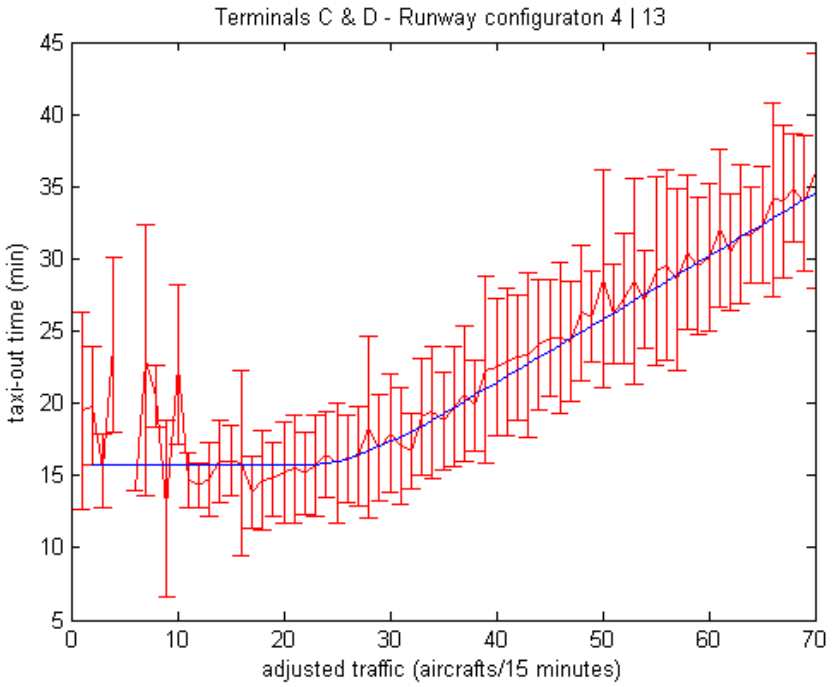

Figure A. XXX: Mean, standard error, and fit function for flights from Terminals C&D when the airport operates with runway configuration 4|13. Data from July 1, 2013 to August 30, 2013

II. APPENDIX B: SATURATION PLOTS

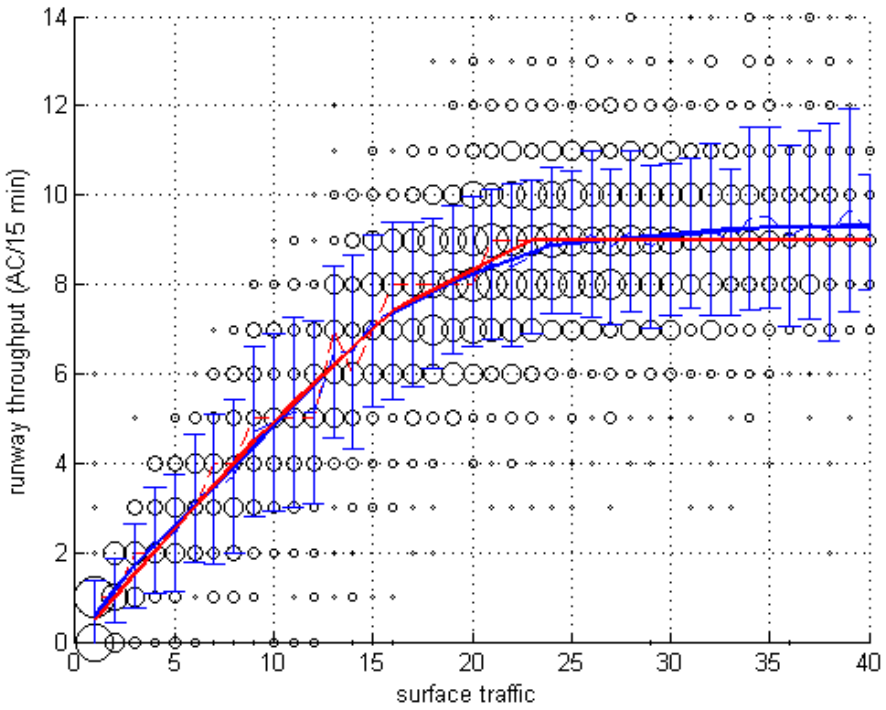

Figure B. I: Saturation plot for runway configuration 31|4 under VMC conditions

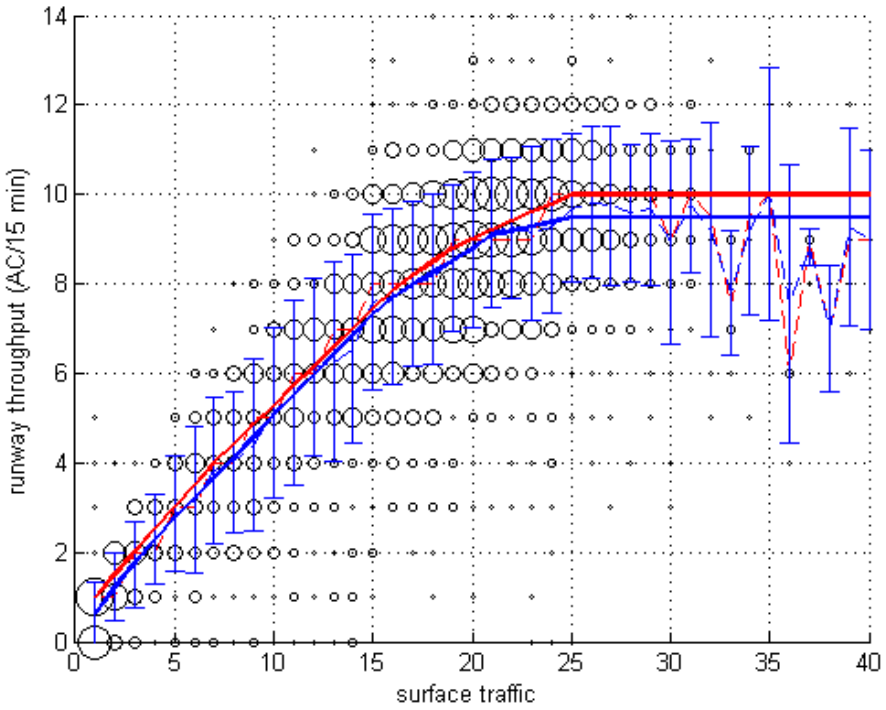

Figure B. II: Saturation plot for runway configuration 22|31 under VMC conditions

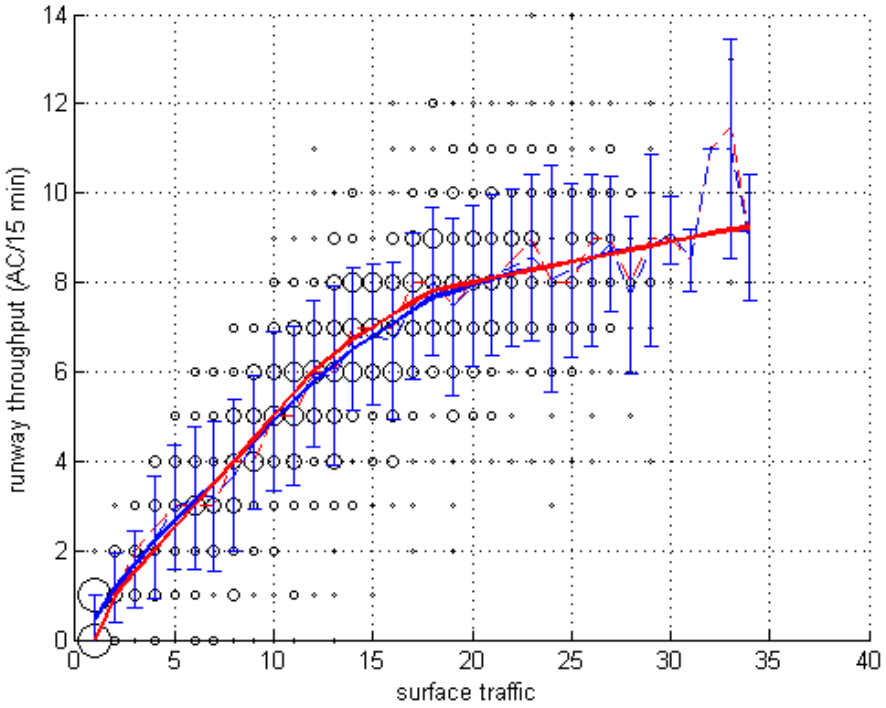

Figure B. III: Saturation plot for runway configuration 31|31 under VMC conditions

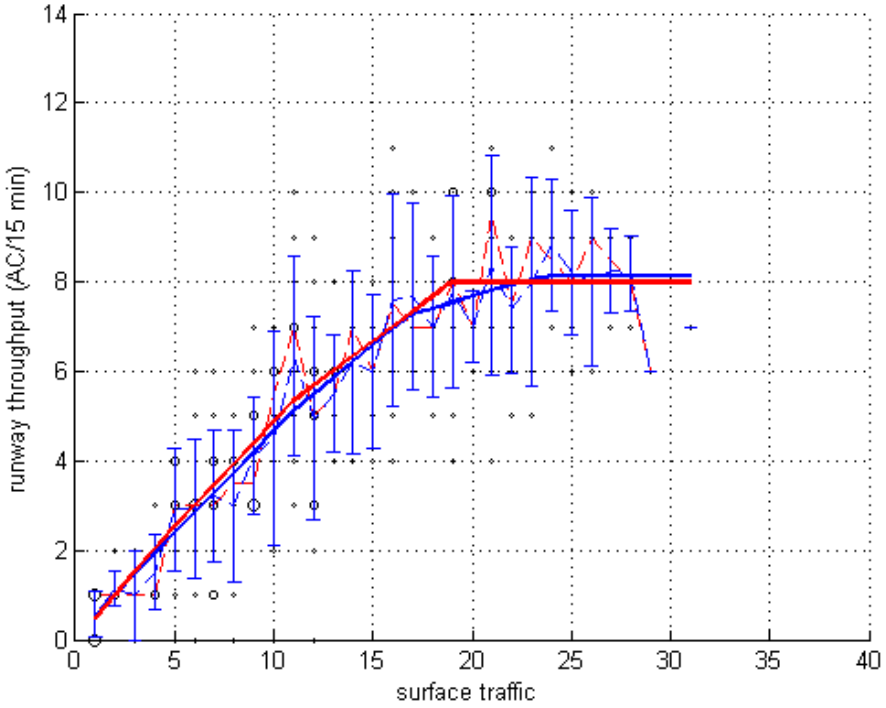

Figure B. IV: Saturation plot for runway configuration 4|31 under VMC conditions

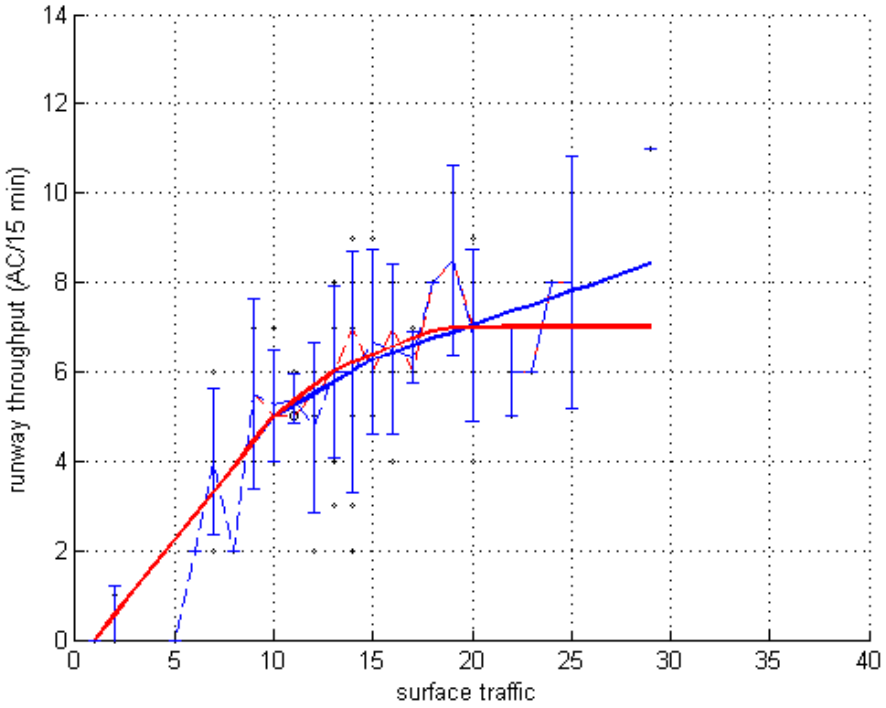

Figure B. V: Saturation plot for runway configuration 13|13 under VMC conditions

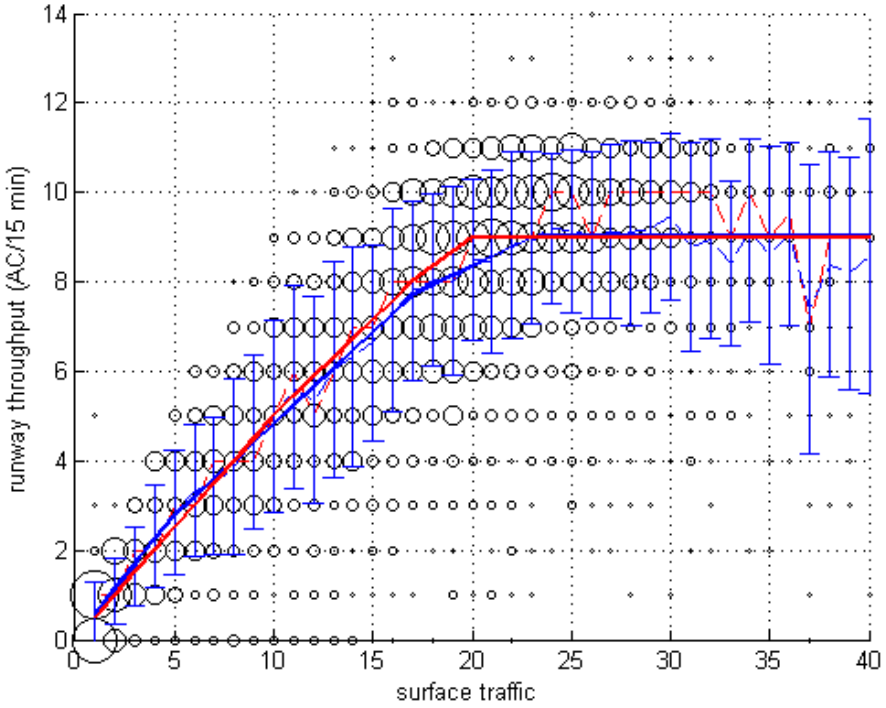

Figure B. VI: Saturation plot for runway configuration 22|13 under VMC conditions

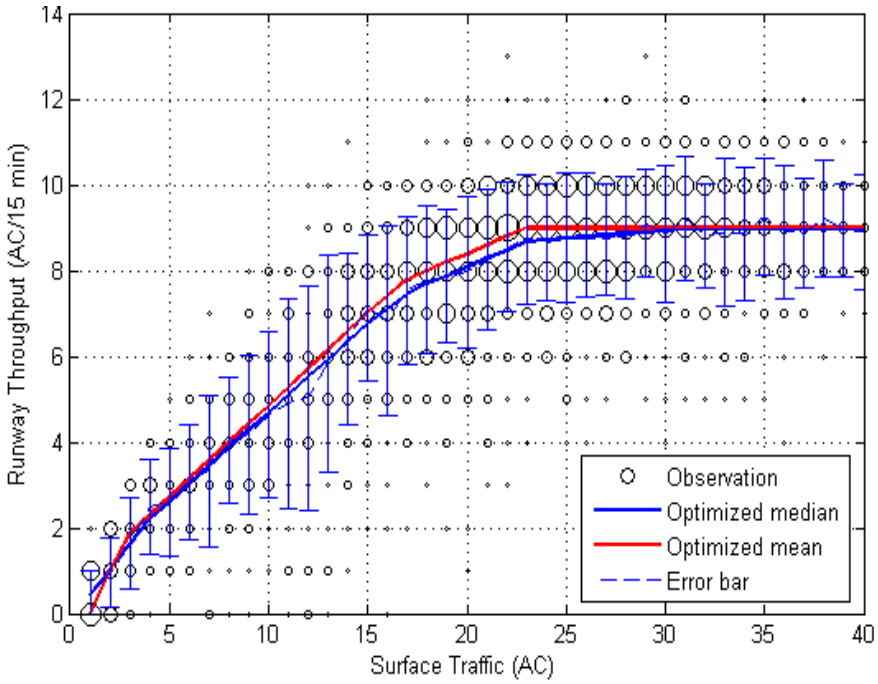

Figure B. VII: Saturation plot for runway configuration 4|13 under VMC conditions

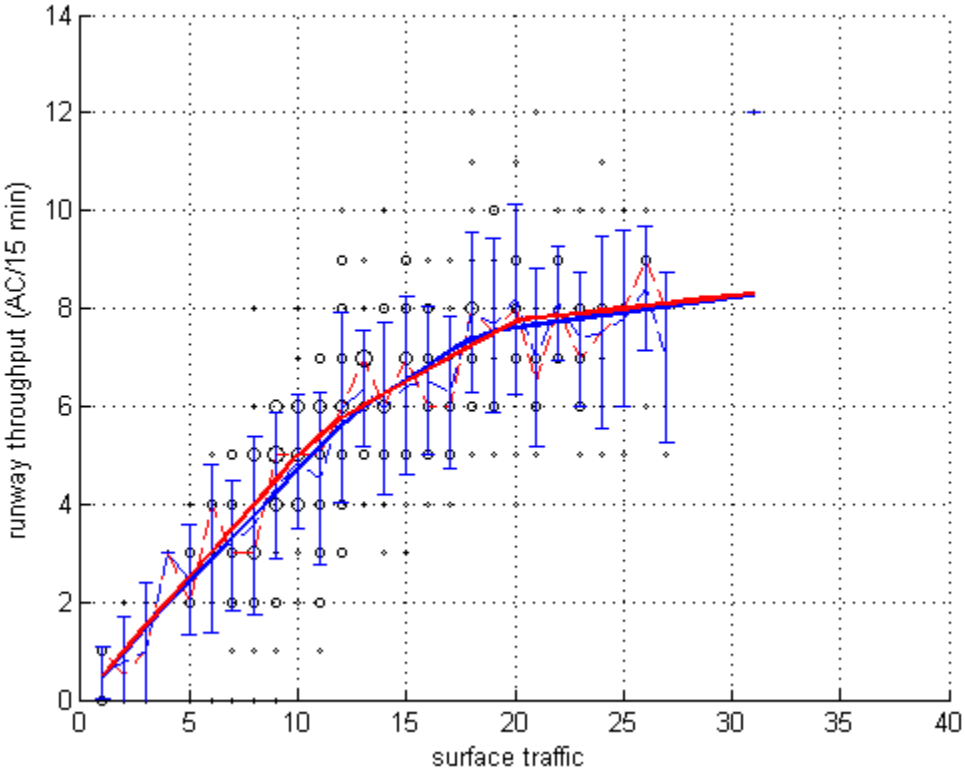

Figure B. VIII: Saturation plot for runway configuration 4|4 under VMC conditions

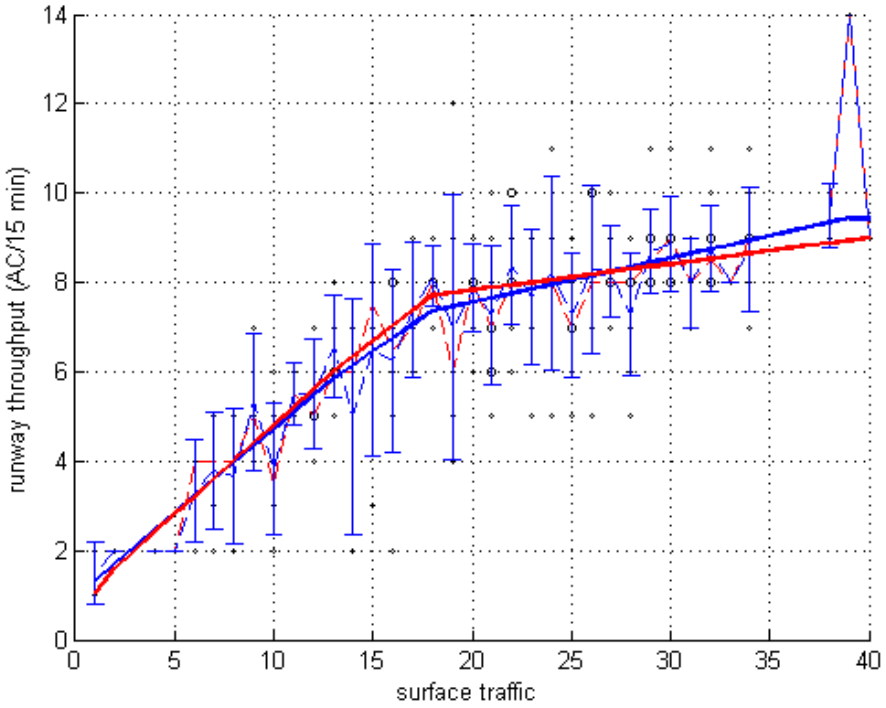

Figure B. IX: Saturation plot for runway configuration 31|4 under IMC conditions

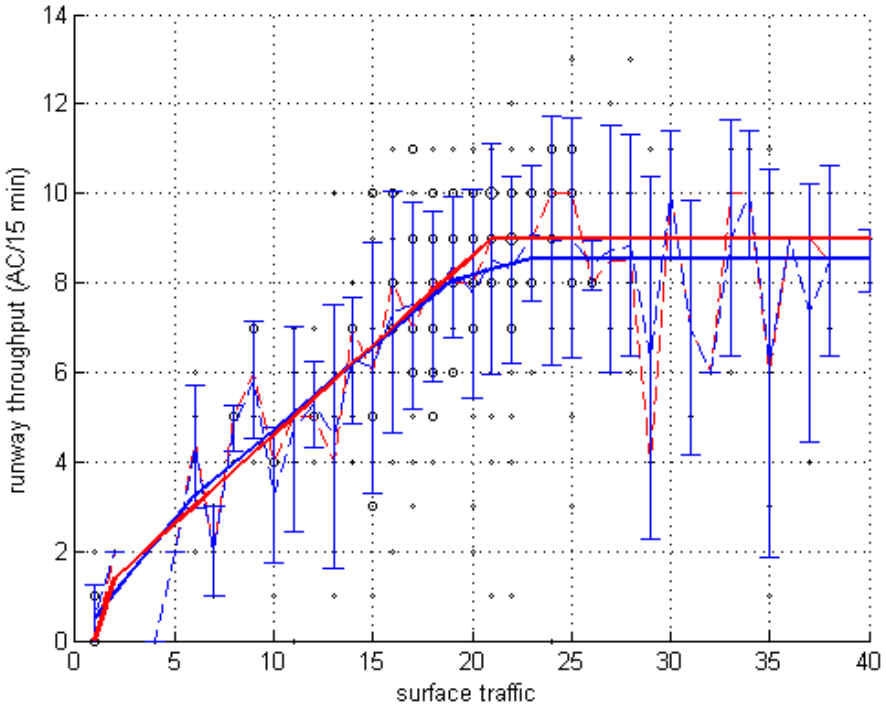

Figure B. X: Saturation plot for runway configuration 22|31 under IMC conditions

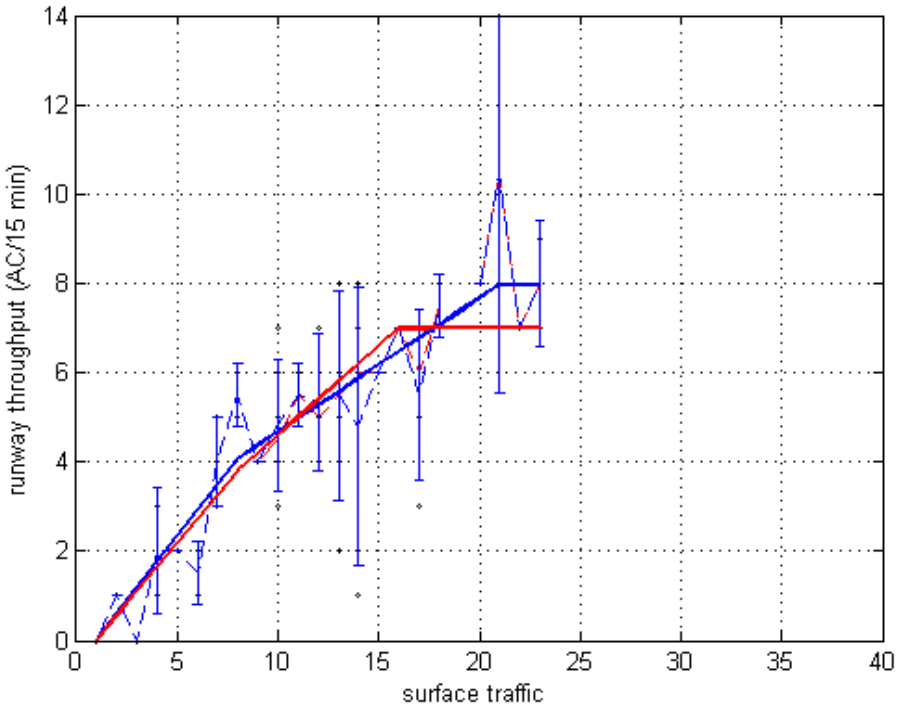

Figure B. XI: Saturation plot for runway configuration 31|31 under IMC conditions

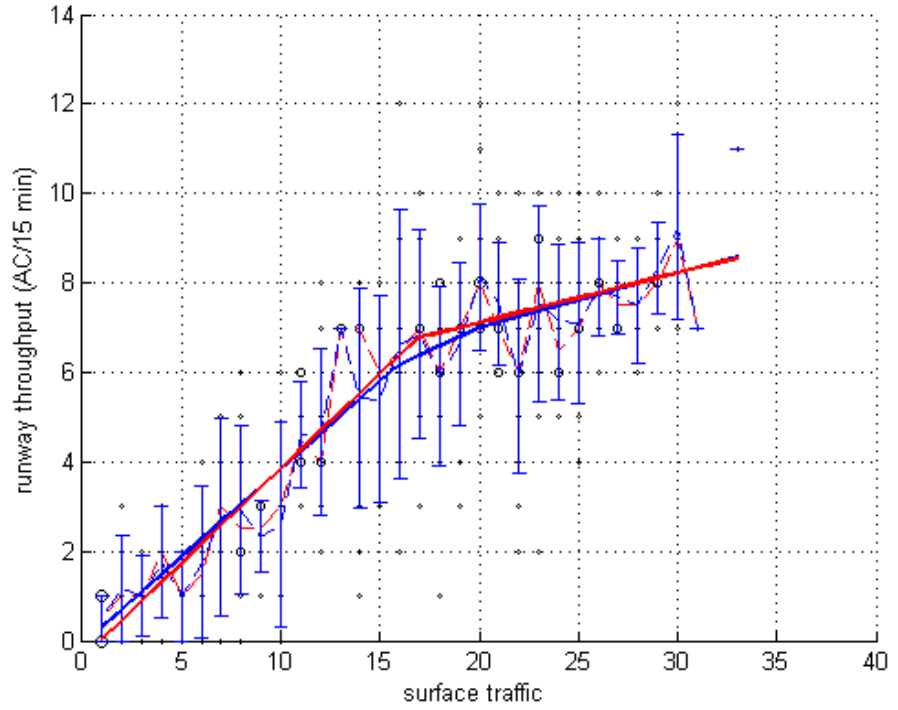

Figure B. XII. Saturation plot for runway configuration 4|31 under IMC conditions

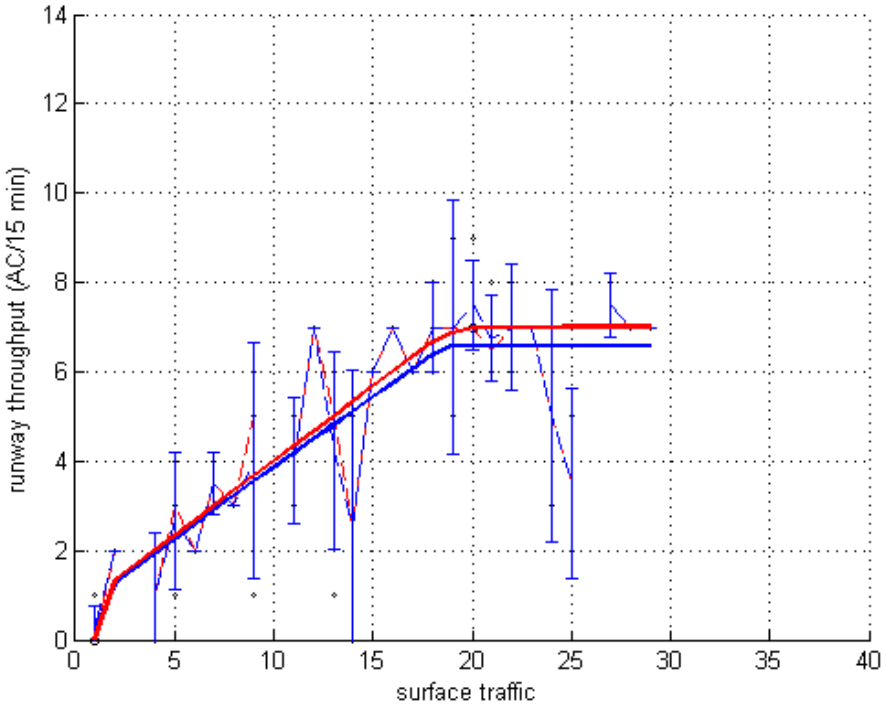

Figure B. XIII: Saturation plot for runway configuration 13|13 under IMC conditions

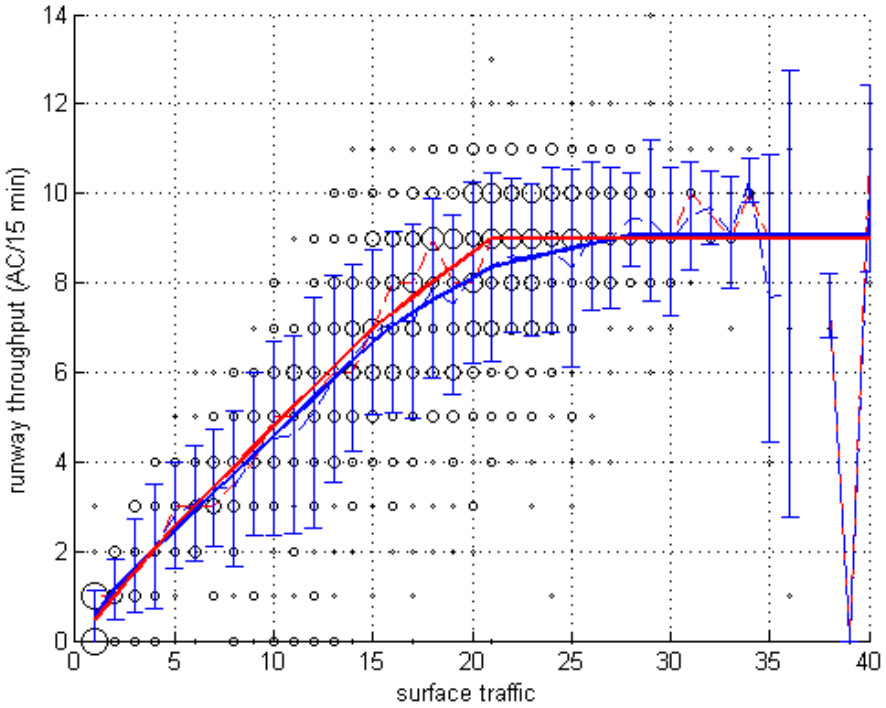

Figure B. XIV: Saturation plot for runway configuration 22|13 under IMC conditions

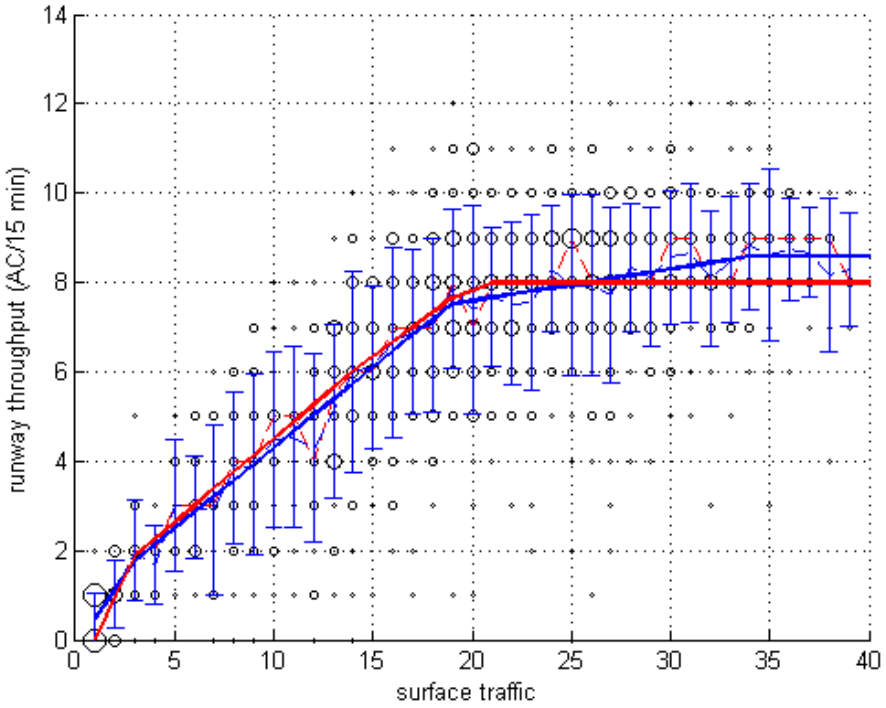

Figure B. XV: Saturation plot for runway configuration 4|13 under IMC conditions

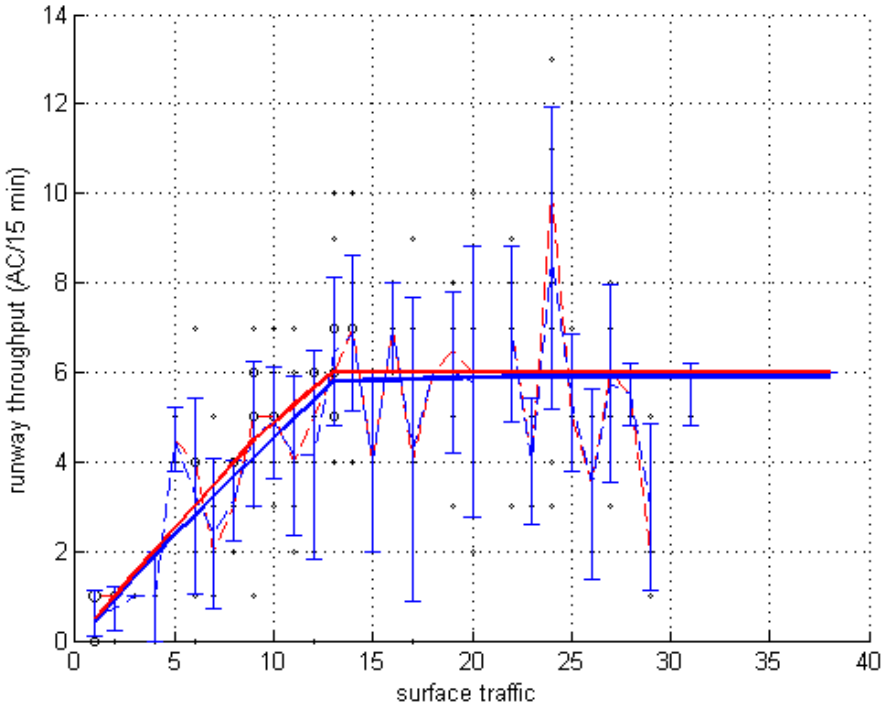

Figure B. XVI: Saturation plot for runway configuration 4|4 under IMC conditions

REFERENCES

Burgain, Pierrick, Eric Feron, and John-Paul Clarke. 2008. “Collaborative Virtual Queue: Fair Management of Congested Departure Operations and Benefit Analysis.” *arXiv Preprint arXiv:0807.0661*. http://arxiv.org/abs/0807.0661.

Carr, Francis, Antony Evans, Eric Feron, and J.-P. Clarke. 2002. “Software Tools to Support Research on Airport Departure Planning.” In *Digital Avionics Systems Conference, 2002. Proceedings. The 21st*, 1:1D3–1. IEEE. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1067900.

Clewlow, R. 2010. *A Prediction Model for Aircraft Surface Delay*. MIT.

DeLaura, Rich, Michael Robinson, Russell Todd, and Kirk MacKenzie. 2008. “Evaluation of Weather Impact Models in Departure Management Decision Support: Operational Performance of the Route Availability Planning Tool (RAPT) Prototype.” In *13th Conference on Aviation, Range, and Aerospace Meteorology, AMS, New Orleans, LA*. https://ams.confex.com/ams/pdfpapers/132880.pdf.

De Neufville, Richard, Amedeo R. Odoni, Peter Belobaba, and Tom Reynolds. 2013. *Airport Systems : Planning, Design, and Management, Second Edition / Richard de Neufville, Amedeo R. Odoni ; with Contributions by Peter Belobaba and Tom Reynolds*. New York : McGraw-Hill Education, c2013 (Norwood, Mass. : Books24x7.com [generator]).

Federal Aviation Administration. “Aviation System Performance Metrics.” *2014*.

Feron, Eric R., R. John Hansman, Amadeo R. Odoni, R. B. Cots, B. Delcaire, W. D. Hall, H. R. Idris, A. Muharremoglu, and N. Pujet. 1997. “The Departure Planner: A Conceptual Discussion,” December. http://dspace.mit.edu/handle/1721.1/34944.

Hall, William D., and Alicia Fernandes. 2013. “Key Performance Issues in Surface Collaborative Decision Making.” In *Digital Avionics Systems Conference (DASC), 2013 IEEE/AIAA 32nd*, 1B1–1. IEEE. http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6712507.

Idris, Husni, John-Paul Clarke, Rani Bhuva, and Laura Kang. 2001. “Queuing Model for Taxi-out Time Estimation.” http://dspace.mit.edu/handle/1721.1/37322.

Jacquillat, Alexandre, and Amedeo R. Odoni. 2015. “Endogenous Control of Arrival and Departure Service Rates in Dynamic and Stochastic Queuing Models of Airport Congestion.” *Transportation Research Part E: Logistics and Transportation Review* 73: 133–51.

Joint Economic Committee. 2008. *Flight Delays Cost Passengers, Airlines, and the U.S. Economy Billions*.

Khadilkar, Harshad. 2011. *Analysis and Modeling of Airport Surface Operations / by Harshad Khadilkar*. c2011. http://libproxy.mit.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=cat00916a&AN=mit.002002779&site=eds-live.

Kim, Sang Hyun, and Eric Feron. 2014. “Impact of Gate Assignment on Departure Metering.” *IEEE Transactions on Intelligent Transportation Systems* 15 (2): 699–709. doi:10.1109/TITS.2013.2285499.

Office of Aviation Policy and Plans, Federal Aviation Administration. 2002. *Documentation for the Aviation System Performance Metrics (ASPM),*.

Pujet, Nicolas. 1999. “Modeling and Control of the Departure Process of Congested Airports.” Thesis, Massachusetts Institute of Technology. http://dspace.mit.edu/handle/1721.1/9363.

Pujet, Nicolas, Bertrand Delcaire, and Eric Feron. 2003. “Input-Output Modeling and Control of the Departure Process of Congested Airports.” http://arc.aiaa.org/doi/pdf/10.2514/6.1999-4299.

Pyrgiotis, Nikolaos. 2012. “A Stochastic and Dynamic Model of Delay Propagation within an Airport Network for Policy Analysis.” Massachusetts Institute of Technology. http://dspace.mit.edu/handle/1721.1/71452.

Ravizza, Stefan, Jun Chen, Jason Atkin, Edmund Burke, and Paul Stewart. 2013. “The Trade-off between Taxi Time and Fuel Consumption in Airport Ground Movement.” *Public Transport (1866749X)* 5 (1/2): 25.

Sandberg, M., I Simaiakis, H. Balakrishnan, T.G. Reynolds, and R.J. Hansman. 2014. “A Decision Support Tool for the Pushback Rate Control of Airport Departures.” *IEEE Transactions on Human-Machine Systems* 44 (3): 416–21. doi:10.1109/THMS.2014.2305906.

Shumsky, Robert Arthur. 1995. “Dynamic Statistical Models for the Prediction of Aircraft Take-off Times.” Thesis, Massachusetts Institute of Technology. http://dspace.mit.edu/handle/1721.1/10592.

Simaiakis, Ioannis. 2009. “Modeling and Control of Airport Departure Processes for Emissions Reduction.” Massachusetts Institute of Technology. http://dspace.mit.edu/handle/1721.1/58289.

Simaiakis, Ioannis. “Analysis, Modeling and Control of the Airport Departure Process.” Massachusetts Institute of Technology. http://dspace.mit.edu/handle/1721.1/79342.

Simaiakis, Ioannis, and Hamsa Balakrishnan. 2009. “Queuing Models of Airport Departure Processes for Emissions Reduction.” In *AIAA Guidance, Navigation and Control Conference and Exhibit*. http://arc.aiaa.org/doi/pdf/10.2514/6.2009-5650.

Simaiakis, Ioannis, and Hamsa Balakrishnan. 2010. “Impact of Congestion on Taxi Times, Fuel Burn, and Emissions at Major Airports.” *Transportation Research Record: Journal of the Transportation Research Board* 2184 (-1): 22–30. doi:10.3141/2184-03.

Simaiakis, Ioannis, Hasmsa Balakrishnan, Harshad Khadilkar, Tom G. Reynolds, R. John Hansman, Brendan Reilly, and Steve Urlass. 2011. “Demonstration of Reduced Airport Congestion through Pushback Rate Control.” http://18.7.29.232/handle/1721.1/60882.