Project 016 Airport Surface Movement Optimization

Massachusetts Institute of Technology
Massachusetts Institute of Technology Lincoln Laboratory

Project Lead Investigator

Hamsa Balakrishnan
Associate Professor
Aeronautics and Astronautics
Massachusetts Institute of Technology
77 Massachusetts Ave, 33-328, Cambridge, MA 02139
(617) 253 6101
hamsa@mit.edu

University Participants

MIT Lincoln Laboratory

- Tom Reynolds
- FAA Award Number: IA # DTFAWA-11-X-80007, FAA Task # 18
- Period of Performance: 1 April 2015 to 31 May 2016
- Task(s):
    1. Site adaptation to different airports
    2. Analysis of the impacts of departure metering in different operating environments
    3. Identification and evaluation of barriers to implementation
    4. Design of implementation protocols and field-testing at selection of study airports
    5. Coordination with Advances Surface Management Programs

MIT

- Hamsa Balakrishnan
- Lincoln subcontract number: Purchase Order 7000332744
- Period of Performance: 1 September 2015 to 31 August 2016
    6. See tasking above

Project Funding Level

$100,000

Investigation Team

Hamsa Balakrishnan (PI -- MIT), Tom Reynolds (PI – MIT Lincoln Laboratory), Patrick McFarlane (graduate student researcher, MIT), Sandeep Badrinath (graduate student researcher, MIT), Melanie Sandberg (associate staff, MIT Lincoln Laboratory).

Project Overview

Aircraft taxiing on the surface contribute significantly to fuel burn and emissions at airports. The quantities of fuel burned as well as different pollutants such as carbon dioxide, hydrocarbons, nitrogen oxides, sulfur oxides and particulate matter are proportional to the taxi times of aircraft, in combination with other factors such as the throttle settings, number of engines that are powered, and pilot and airline decisions regarding engine shutdowns during delays. Domestic flights in the United States in 2008 emitted about 6.6 million US tons of CO_2, 49,000 US tons of CO, 8,800 US tons of NOx, and 4,400 US tons of HC taxiing out for takeoff; almost half of these emissions occurred at the 20 most congested airports in the country. Similar trends have been noted at major airports in Europe, where it is estimated that aircraft spend 10-30% of their flight time taxiing, and that a short/medium range A320 expends as much as 5-10% of its fuel on the ground. The purpose of the Airport Surface Movement Optimization study is to show that a significant portion of these impacts can be reduced through "technologically-lightweight" operational measures to limit airport surface congestion.

A simple airport congestion control strategy would be a pushback policy aimed at reducing congestion on the ground that would consider the situation on the airport surface (also called the state). The N-control strategy is one such approach, and was first considered in the Departure Planner project. Several variants of this policy have been studied since in literature. The policy is effectively a simple threshold heuristic: if the total number of departing aircraft on the ground exceeds a certain threshold, further pushbacks are stopped until the number of aircraft on the ground drops below the threshold. In our early analysis we discovered that this form of discrete, on-off control strategy was difficult to implement in practice, and could also be overly reactive, potentially leading to instability. By contrast, the pushback rate control strategy that we have developed and tested at Boston Logan International airport (BOS) does not stop pushbacks once the surface is in a congested state, instead it regulates the rate at which aircraft pushback from their gates during high departure demand periods so that the airport does not reach undesirably high congested states. This document summarizes the Phase 2 efforts, including site selection criteria and developing techniques for characterizing airport surface operations, in order to enable the adaptation of a given congestion management approach to different airports, and the comprehensive evaluation of implementations.

As part of this project, MIT undertook an initial assessment of the applicability and potential benefits of "light-weight" airport-wide surface management control concepts involving minimal levels of automation to complement other Federal Aviation Administration (FAA) surface congestion management programs. It involved defining and modeling surface management control schemes, implementing them in a field demonstration at Boston Logan International airport, and evaluating performance in terms of impacts on taxi time, fuel burn and environmental emissions. During 15 four-hour tests conducted during the summers of 2010 and 2011, fuel use was reduced by an estimated 23-25 US tons (6,600-7,300 US gallons), while carbon dioxide emissions were reduced by an estimated 71-79 US tons. These savings were achieved with average gate-holds of just 4.7 min, and savings of 114-128 lb of fuel per gate-held flight. In addition to these savings achieved during field trials, many important lessons were learned regarding operational implementation of surface management techniques in both nominal and off-nominal conditions.

Most prior research (including this project to date) has focused on demonstrations of a proposed congestion management approach at a particular airport, and not on the adaptation of a particular approach to a range of airport operating environments. The current focus of this project addresses the challenges involved with adapting any class of surface congestion management approaches to different airports. Data and case studies from New York's LaGuardia Airport and Philadelphia International Airport are used to illustrate the diversity in operating environments. In particular, the MIT team has developed techniques for characterizing airport surface operations using site surveys and operational data. These characterizations are used for the adaptation of a given congestion management approach to different airports, and for the comprehensive evaluation of implementations.

Integration of Departure Metering Concepts into Surface Capabilities

Objective(s)

The objectives of this project are to conduct an initial assessment of the applicability and potential benefits of relatively easy-to-implement airport-wide surface management control concepts involving minimal levels of automation and procedural modifications, to complement other FAA surface congestion management programs. Phase 1 involved defining and modeling surface management control schemes, implementing them in a field demonstration at Boston Logan International airport, and evaluating performance in terms of impacts on taxi time, fuel burn and environmental emissions. Phase 2, conducted during FY14 and 15, is exploring adaptation of the approach to other airport locations with very different operating characteristics to help understand and inform requirements for more general deployment in future FAA decision support tools.

Research Approach

Framework for adapting approaches to different operating environments

This study has identified the overall process for designing a congestion management approach illustrated in Figure 1. The main steps involved in this process are: (1) Airport Selection, where an airport with surface congestion problems are identified; (2) Airport Characterization, where the details of the operation relevant to surface congestion management at an airport are identified; (3) Algorithm Development, where specific surface congestion management approaches are created; (4) Implementation Design, where the protocols of the execution of the algorithms are developed for the airport; and (5) Operational Testing and Performance Evaluation, where the approach is tested and evaluated in the operational setting.

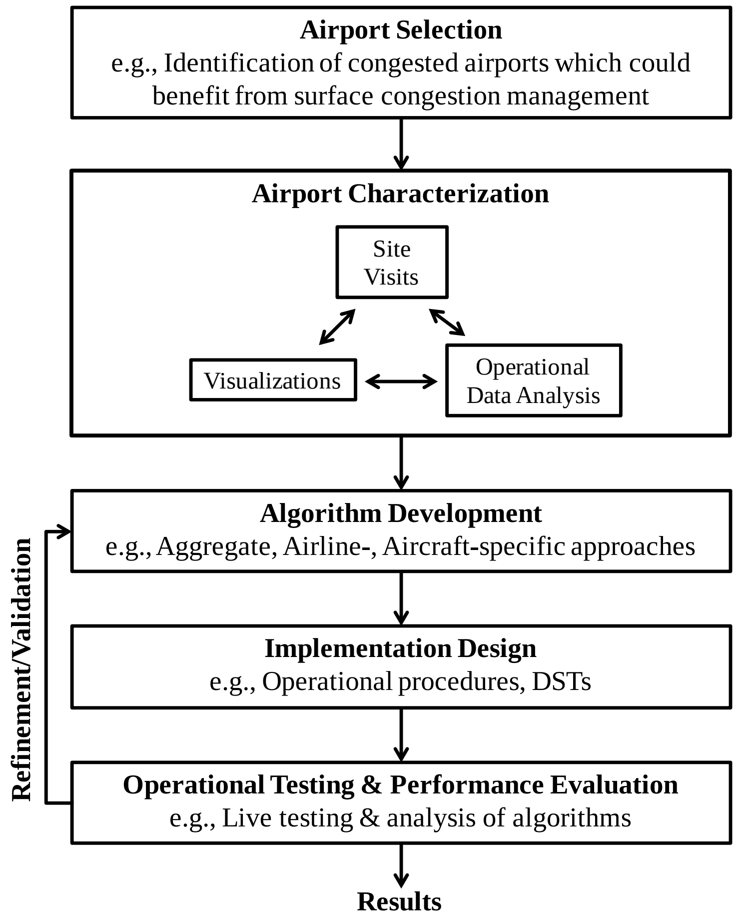

Figure 1. Overall design process for a congestion management approach.

The airport selection step resulted in an analysis focus on LGA and PHL airports during Phase 2 activities.

Framework for incorporating departure metering into Surface CDM

Figure 2 presents a schematic of the S-CDM concept of operations, with a particular emphasis on the data elements required, and the outputs.

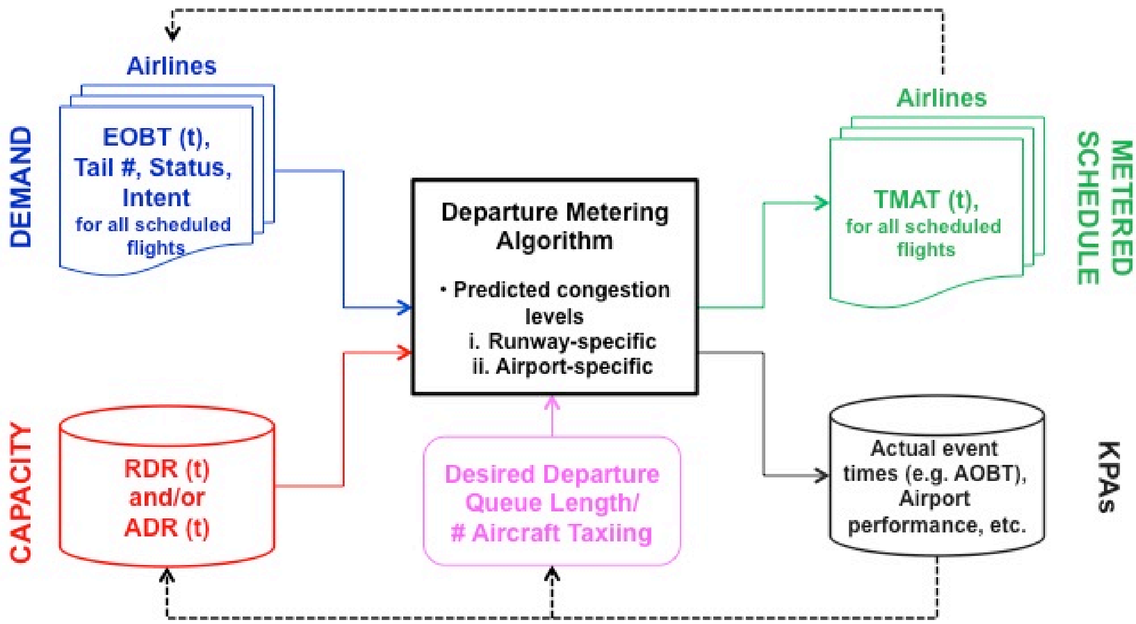

Figure 2. Simplified schematic of S-CDM ConOps, assuming availability of perfect data.

In prior work, research team has developed and validated data-driven algorithms for determining the appropriate values of the desired departure queue length and the number of aircraft taxiing (the quantities shown in magenta), given tactical (i.e., 15-minutes ahead) estimates of the demand and operating conditions. The algorithms account for the variability/uncertainty in throughput (RDR/ADR) over the next 15-minute period.

Departure metering using Pushback Rate Control strategies

Two classes of Pushback Rate Control (PRC) strategies were developed and analyzed: (1) An N-Control policy, and (2) a more advanced dynamic programming based policy. N-control aims to maintain departing traffic around an acceptable value N_{ctrl} based on empirical analysis of the relationship between departure throughput and active traffic, while dynamic programming seeks to minimize a cost function that penalizes both long queues and runway starvation by calculating the probability of the airport being in a certain state at some future time. Both approaches use predictions of the departure throughput in a time-window of interest by considering the runway configuration, weather conditions, and arrival demand using regression trees. Information on downstream restrictions is used for more accurate predictions of the operational throughput of an airport by leveraging the Route Availability Planning Tool (RAPT), an automated decision support tool that identifies departure routes that will be impacted by convective weather. Archived RAPT data can be used to predict the impact of route availability on the capacity of the airport. Regression trees are used to predict the departure throughput of LGA in each 15-minute interval as a function of arrival rate and a “RAPT value”, which is used to measure the level of route blockage in the departure airspace.

N-Control based PRC: The N-Control policy considers the runway configuration, meteorological conditions, demand and RAPT forecast in a 15-minute period in order to determine the suggested pushback rate for that time period. A schematic of this process is shown in Figure 3.

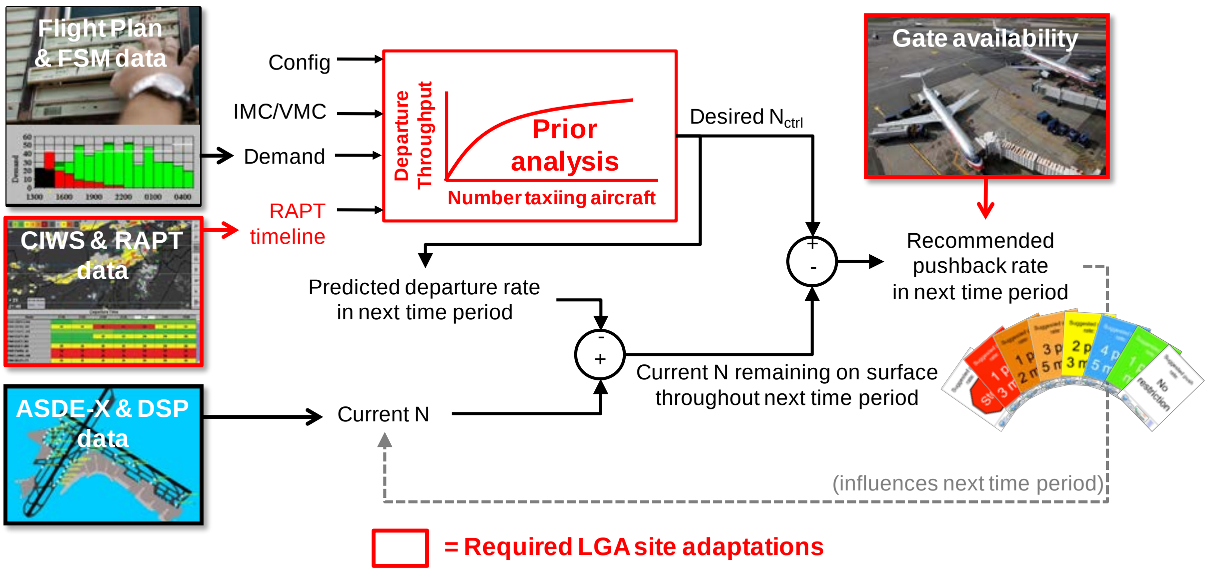

Figure 3. Adaptation of N-control to derive a PRC policy for LGA.

Dynamic Programming based PRC: Dynamic programming can be used to determine PRC policies that control the departure pushback rate by minimizing a cost function that penalizes long runway queues and low runway utilization. For each time-window (say, 15 min) and given the current state of the system, a queuing model is used to predict the state of the system at the end of that time-window. A prediction of the departure throughput of the airport for that time-window (which in turn depends on the arrival rate of the airport) is required in order to predict the state of the airport. The dynamic program then determines an optimal pushback rate for the duration of the time-window. The dynamic programming policy has the benefit that it accommodates probabilistic predictions of throughput. Figure 4 illustrates such a policy derived for LGA.

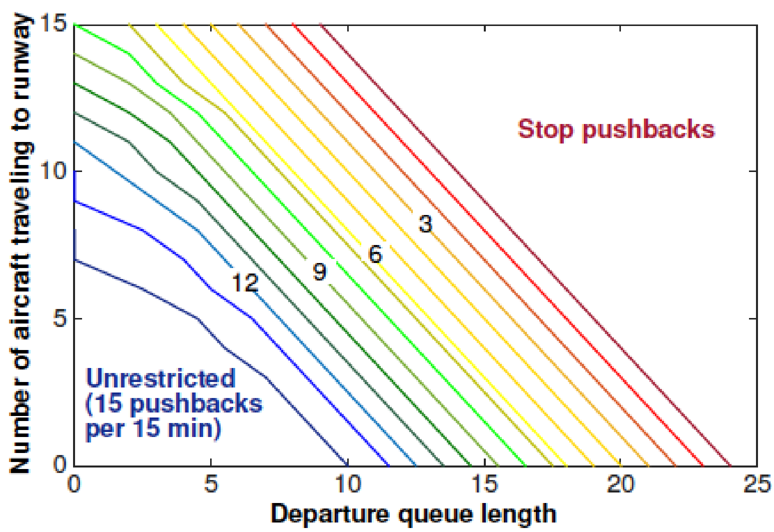

Figure 4. Dynamic programming based pushback rate policy for a time-window of 15 min.

The dynamic programming policy contrasts with the N-control policy in the following manner. N-control uses a simple equation to maintain departure surface traffic at a predetermined level based on empirical data. Dynamic programming models the runway service time to get a probability distribution of the state of the airport at some point in the future. With this, and a cost of queuing and runway utilization function, dynamic programming finds the departure pushback rate that minimizes costs. While N-control predicts that the state of the airport surface will evolve in a certain manner, dynamic programming considers all of the potential states of the airport and the departure pushback rate accounts for the uncertainty in the evolution of the state of the airport. From this perspective, the dynamic programming approach, while more complex than the N-Control based approach, is a more robust policy. In addition, while N-control is only in effect during times of significant congestion ($N > N_{ctrl}$), the dynamic programming policy is always in effect. As a result, the dynamic programming based approach remedies even smaller (temporary) imbalances between throughput and demand, and results in greater taxi-out time reductions.

Time windows and time horizons

The PRC policies, N-control and dynamic programming, both generate a pushback rate for departures valid for a given time window. Historically, this time window has been set to 15 minutes, but varying this time window can lead to advantages and disadvantages. Also, a pushback rate can be calculated for earlier time windows in the future by changing the time horizon of the policy. Changing the time window or time horizon allows an airport or airline to tailor the PRC policy to specific needs and requirements. Consequently, stakeholders need to be aware of how varying these parameter values can affect the performance of PRC policies. Figure 5 illustrates the difference between time windows and time horizons.

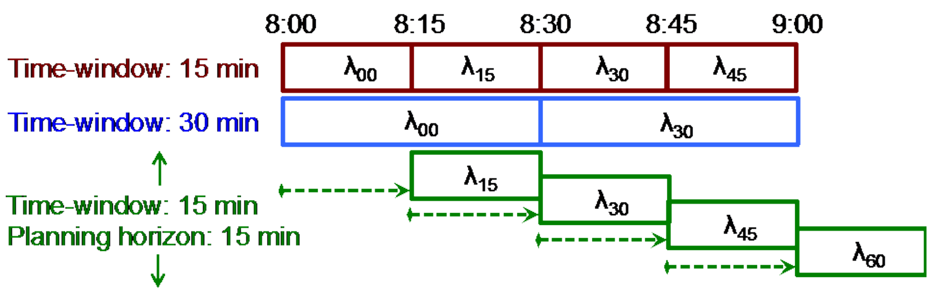

Figure 4. Visualization of time windows and time horizons.

PRC policies calculate a pushback rate for departing aircraft that is valid for a certain time window. This time window impacts some performance characteristics of the policy. The length of this time window is a tradeoff between accuracy, ease of implementation, and value added to operators (airlines and air traffic controllers). For accuracy, the policies become less accurate as the length of the time window increases. Because the PRC policies calculate the pushback rate at the beginning of the time window based on the state of the airport surface, a longer time window means that the pushback rate is valid (recommended) for a longer period of time. As time gets farther away from the beginning of the window, the state of the airport surface changes relative to what was expected at the start of the time window, resulting in a decrease in metering performance.

The time horizon is the number of time windows before a given time window that a departure rate is calculated. The time horizon length variation has similar tradeoffs to those associated with time window length variation. Extending time windows or horizons reduces operator workload resulting from the policy, but prediction and data accuracy also decrease. Shortening time windows or horizons increases accuracy, but operator workload also increases due to short planning horizons.

Impact of uncertainty

The uncertainty associated with the throughput prediction increases if either the time window increases (say, to 30 min instead of 15 min), or if the time horizon increases (for example, planning for the time-windows that begin 15- or 30-min later instead of just at the current time). On the demand-side, arrival rate predictions may not be accurate, which will in turn increase the uncertainty associated with the predicted departure throughput. Similarly, departure demand may vary and aircraft may not push back at the times recommended by the pushback rate.

Departure demand (EOBT) uncertainty

Both N-control and dynamic programming rely on departure demand being ready for pushback in a given time window, as expected when determining the pushback rate. In the S-CDM context, this translates to an accurate knowledge of the Earliest Off-Block Times (EOBTs). If the available departures are less than the pushback rate, the departure surface traffic may fall below acceptable levels. If there are many available departures during a time of low congestion, those departures may cause congestion in the future. As such, PRC policies need to handle variability in departure demand. Such variability is modeled in simulations of LGA operations by assuming that the perturbations in the EOBT are drawn from a normal distribution with a mean of the published EOBT and a standard deviation of 3.5 minutes. The assumption of a standard deviation of 3.5 minutes ensures that two standard deviations from the expected EOBT approximately encompasses a 15-minute time window.

Arrival rate uncertainty

The number of arrivals landing at an airport in a given time window is also uncertain. The scheduled arrival time of an aircraft changes due to many of the reasons that affect the scheduled departure time, including weather and congestion. Instead of scheduled arrival times, a tool called the Flight Status Monitor (FSM) forecasts the number of arrivals that will be ready to land for time windows into the future. The FAA's Airport Arrival Demand Chart (AADC) is used as a surrogate for the arrival rate predictions from the FSM. Figure 5 shows an example of the AADC predictions for LGA.

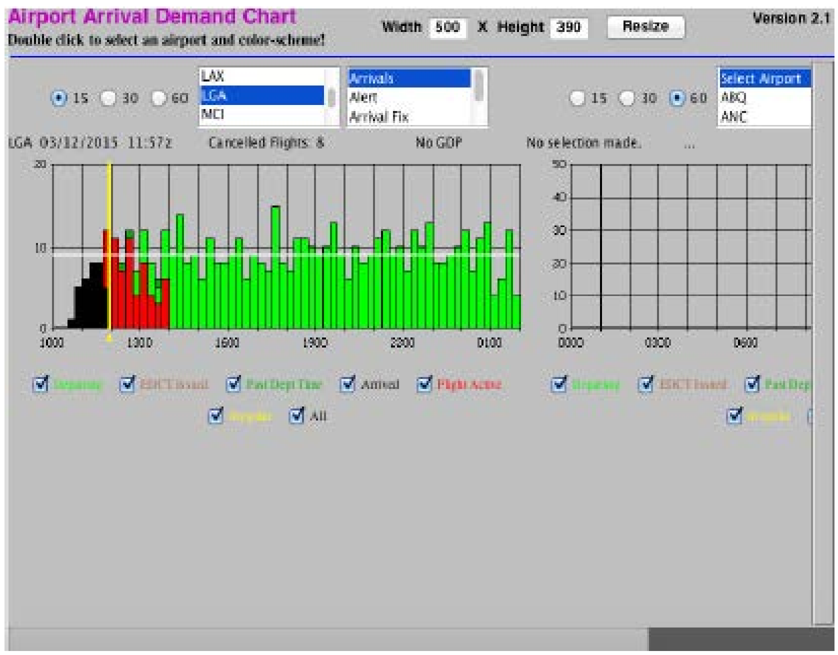

Figure 5. Airport Arrival Demand Chart for LGA on 3/12/2015. The bars indicate how many aircraft will be approaching LGA for landing during each 15-minute time window.

With this information, the predicted arrivals for the next 15, 30, and 60 minutes can be extracted. If the predicted number of arrivals is above the capacity, the predicted arrivals is set to the capacity. The prediction accuracy analysis that follows establishes the perturbations for the variable arrival rate analysis. The AADC data was gathered over a four-month period from January-April 2015. Figure 6 shows the distribution of the difference between AADC arrival predictions and the actual arrivals for 15-minute windows.

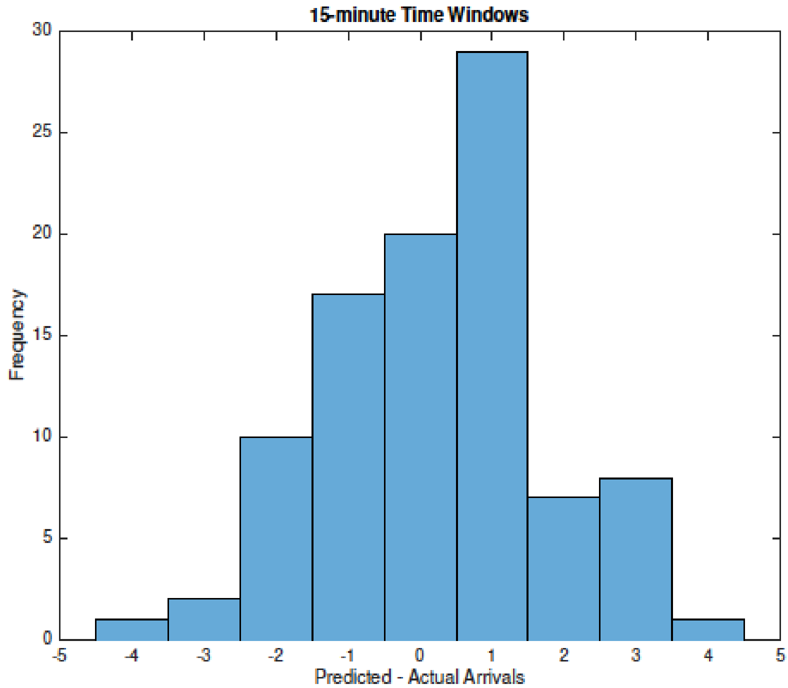

Figure 6. Predicted minus actual arrivals at LGA for the next 15-minute window.

Simulations of departure metering under uncertainty

Input data

The data required for the simulation are extracted from multiple sources. The ASPM dataset provides flight-specific metrics such as pushback time, wheels off time, and wheels on time. Gate and terminal assignments allow for the calculation of unimpeded taxi-out times, and allow the policy to monitor gate conflicts. The simulations resolve all gate conflicts by clearing a departure to pushback (if its EOBT has passed) as soon as an arrival assigned to that gate has landed. The last dataset contains the weather data, RAPT, described previously. The simulations consider July and August 2013.

Each simulation contains both a baseline case and two metering cases. The baseline case simulates the airport operations by releasing departures from their gates on a First-Come-First-Served (FCFS) basis based on their EOBTs (assumed to be the scheduled departure times with the appropriate perturbation). The metering cases simulate the airport operations using the N-Control and dynamic programming policies. The benefits of departure metering include the taxi-out time reduction, which is the difference between the taxi-out times in the baseline case and metering case. Taxi-out time reduction contrasts with gate holding time, which is the length of time an aircraft is held at a gate beyond the scheduled departure time due to the departure metering policy. Although aircraft still belong to the virtual queue with engines off, occupying the gate for longer periods of time causes more gate conflicts, and extended gate holding times after boarding can inconvenience passengers.

Simulation results

The results of the simulations are presented below. For each departure metering strategy, the plots show the resulting reduction in taxi-out time, namely, the difference between the baseline taxi-out time and the one under the metering policy. As mentioned earlier, comparing the results of the N-control and the dynamic programming based policy simulations highlights some of the differences between the algorithms. The absolute benefits of the dynamic programming policy are therefore much greater than the benefits of the N-control policy. However, this does not immediately indicate that dynamic programming is the better policy. Because the dynamic programming algorithm relies on an infinite horizon solution, the dynamic programming algorithm controls departures at all times, not just during times of congestion like the N-control policy. Therefore, the dynamic programming algorithm meters many more flights than the N-control algorithm. As a result, taxi-out time reductions are expected to be greater under the dynamic programming based policy.

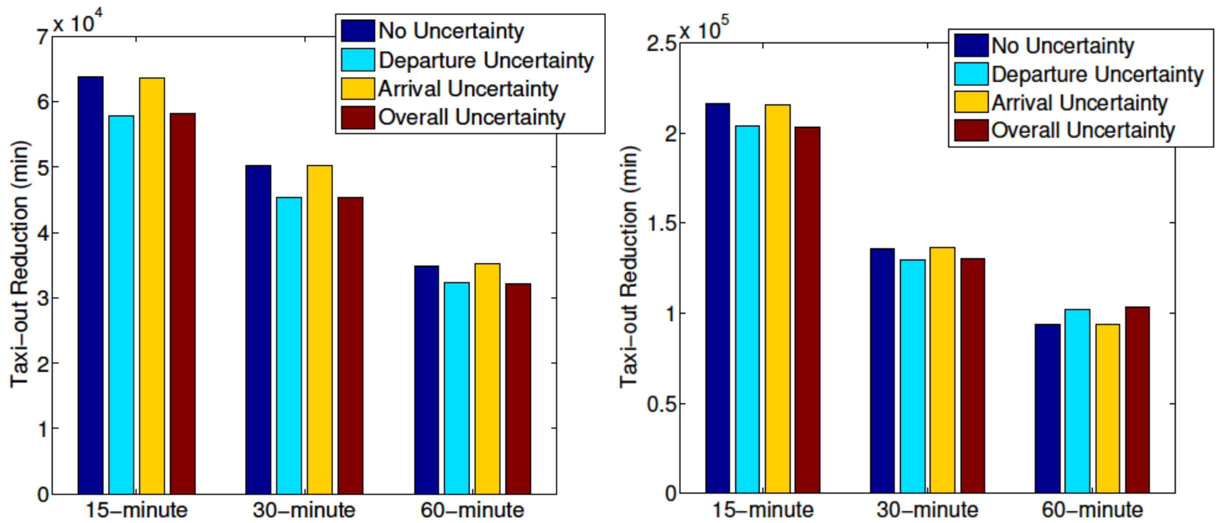

Figure 7. Total taxi-out reduction for the LGA time window simulations, under (left) N-Control, and (b) dynamic programming. Note that the y-axis scales are different in the two plots.

Influence of time window length: Three time windows are investigated: 15-, 30- and 60-minutes. Figure 7 shows the taxi-out time reductions under (left) N-Control, and (right) dynamic programming based PRC. As expected, it shows that the taxi-out time reduction relative to the baseline under the N-control policy (for a time window of 15 min, and no further advance time horizon) is significantly more under the dynamic programming based policy than the N-control policy, for all three values of time window. The results show that the benefits decrease as the time window increases, since a single pushback rate is selected for the entire time window, even if the departure demand fluctuates within it. As a result, the 30-minute time window simulation has 79% of the taxi-out time reduction of the 15-minute time window simulation, while the 60-minute time window simulation has 55% of the taxi-out time reduction of the 15-minute time window simulation. However, even the longer time-windows exhibit considerable benefits, nearly 100,000 min with a 60 min time window over the two month period, using a dynamic programming based pushback rate control policy.

The results show that the EOBT (departure demand) uncertainty has a larger effect on taxi-out reduction than does the arrival rate uncertainty. The overall uncertainty (both EOBT and arrival rate uncertainty) analysis results virtually match the EOBT uncertainty-alone analysis results for all time window lengths. While the 15-minute and 30-minute time window simulations have a decrease in taxi-out reduction due to the EOBT uncertainty, the 60-minute time window simulation has a slight increase in taxi-out reduction. This increase may be due to the spreading out of the departure demand through perturbations, which could result in more flights being metered than initially planned.

Influence of time horizon length: Three time windows were investigated, each with a time window of 15-min. They corresponded to no look ahead (i.e., the pushback rate is determined only for the immediately next 15-minute period), 1 time window look ahead (i.e., a pushback rate was determined for the 15-minute interval that was 15-30 minutes from the current time), and a 3 time window look ahead (i.e., a pushback rate was determined for the 15-minute interval that was 45-60 minutes from the current time). These results, shown in Figure 8, are similar to the results for time window variation shown in Figure 7.

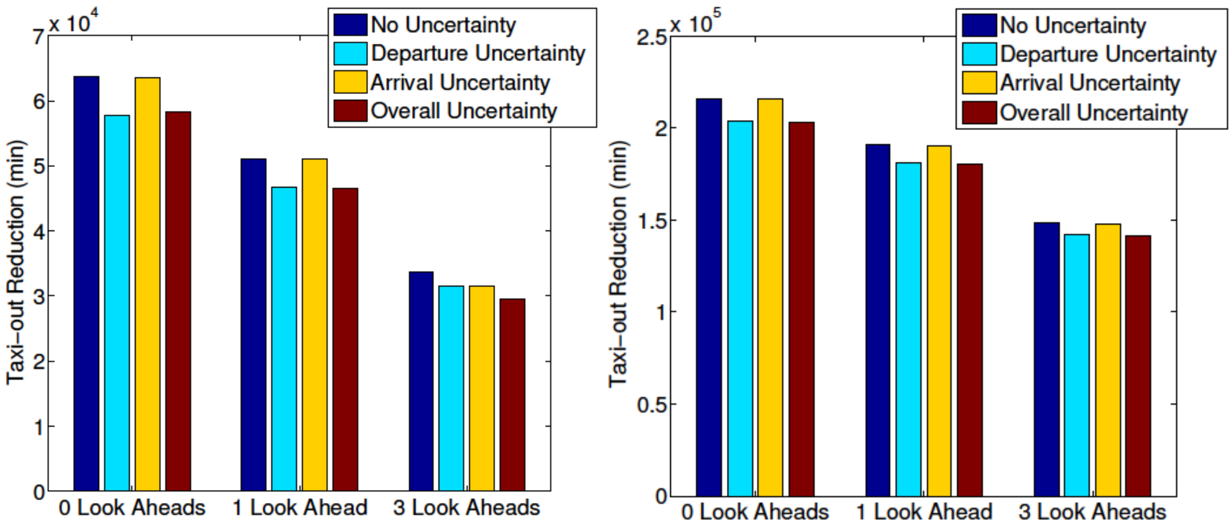

Figure 7. Total taxi-out reduction for the LGA time horizon simulations, under (left) N-Control, and (b) dynamic programming. Note that the y-axis scales are different in the two plots.

The results show that the departure schedule uncertainty has a larger effect on policy benefits than does the arrival rate uncertainty. For both N-control and dynamic programming, the departure schedule uncertainty has less of a negative influence on the results as the time window length increases. In fact, for the 60-minute time window dynamic programming simulation, the departure schedule uncertainty slightly increases the taxi-out reduction. For the time horizon simulations, the influence of each uncertainty remains proportional to the policy benefits with no uncertainty. The arrival rate uncertainty has very little effect on taxi-out reduction for any simulation.

Next Steps and Timeline

Ongoing work consists largely of further refining the proposed framework to support the integration of departure metering with NextGen surface management capabilities. Pursuant to the recent FAA decision to address the RTCA NextGen Integrated Working Group (NIWG) surface team departure metering recommendation with a field demonstration by NASA at Charlotte (CLT) Douglas airport, the research team has been extending the methodologies to CLT, in order to support the ATD-2 demonstration. The team had previously analyzed CLT as part of the site selection task for Phase 2, and also investigated the impacts of the new runway at CLT that became operational in 2011. In particular, discussions with the NASA benefits analysis team lead for the CLT activity has highlighted strong interest in getting a better understanding of departure metering benefits estimates, particularly with the Spot and Runway Departure Advisor (SARDA) algorithm coupled into an S-CDM-compliant architecture, and the framework proposed in this research can enable that. NASA is also concerned with the impact to benefits of uncertainties (e.g., in EOBTs, in departure sequencing times, in gaps in the airspace, etc.) as a function of different look ahead times, similar to the analysis presented in this document.
As the capability is being refined, the team is also coordinating with the TFDM program office and the Surface Operations Office to ensure maximum impact from the N-Control activities. TFDM will be the FAA's primary acquisition program to deliver surface capabilities. The results of these studies were briefed to the TFDM Investment Analysis Team (IAT) to assist with the validation of potential benefits from their departure metering capability.

Milestone(s)

- Site selection [Complete]
- Framework to adapt departure metering concepts to different operating environments [Complete]
- Protocol design for LGA and refinement (multiple) [Complete]
- Initial stakeholder engagement [Complete]
- Modeling and simulation of departure metering at LGA [Complete]
- Modeling and simulation of departure metering at PHL [Complete]

- Engagement/coordination with TFDM benefits analysis [Ongoing]
- Investigation of benefits of incorporation of S-CDM data elements (e.g. EOBT) [Ongoing]
- Identification of most appropriate departure metering algorithm for different airports and operating conditions and levels of uncertainty/ data accuracy (CLT adaptation) [Ongoing]
- Coordination with Advanced Surface Management Programs [Ongoing]

Publications

Peer-reviewed conference publications:
P. McFarlane and H. Balakrishnan. "Optimal Control of Airport Pushbacks in the Presence of Uncertainties," American Control Conference, July 2016.

Outreach Efforts

N/A

Student Involvement

Patrick McFarlane. Modeling and simulation of PHL operations as well as uncertainty analysis. Graduated with Masters degree in February 2016 and currently employed at MITRE.

Sandeep Badrinath, currently a graduate student at MIT. Analysis and control of queuing network models of airport operations.

Plans for Next Period

Task 1: Investigate effect on departure metering algorithms/benefits of incorporation of S-CDM data elements (e.g., EOBT, gate information, etc)

Task 2: Investigate most appropriate departure metering algorithm for different types of airports and operating conditions. Focus on CLT ramp operations.

Task 3: Determine what metrics should be collected to assess and refine strategic and tactical departure metering performance.

Task 4: Coordination with Advanced Surface Management Programs (S-CDM and TFDM)