Airport Characterization for the Adaptation of Surface Congestion Management Approaches

Melanie Sandberg, Tom Reynolds,
Harshad Khadilkar and Hamsa Balakrishnan

Report No. ICAT-2013-01
February 2013

MIT International Center for Air Transportation (ICAT)
Department of Aeronautics & Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139 USA

[Page Intentionally Left Blank]

Airport Characterization for the Adaptation of Surface Congestion Management Approaches[*]

Melanie Sandberg and Tom G. Reynolds

Weather Sensing Group
MIT Lincoln Laboratory
Lexington, MA, USA
{lanie.sandberg, tgr}@ll.mit.edu

Harshad Khadilkar and Hamsa Balakrishnan

Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA, USA
{harshadk, hamsa}@mit.edu

Abstract—**Surface congestion management has received increased attention worldwide, largely due to its potential to mitigate operational inefficiencies and environmental impact. Most prior efforts have 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. This paper illustrates the challenges involved with adapting any class of surface congestion management approaches to different airports. Data and case studies from Boston Logan International Airport, New York's LaGuardia Airport and Philadelphia International Airport are used to illustrate the diversity in operating environments. The paper then proposes techniques for characterizing airport surface operations using site surveys and operational data. Finally, it shows how these characterizations can be used for the adaptation of a given congestion management approach to different airports.**

Keywords-airport congestion control; departure planner; surface manager; deployment; performance characterization

I. INTRODUCTION

Surface congestion is a problem faced by most major airports, and results in increased taxi times, fuel burn and emissions. A recent study of major US airports estimated that Philadelphia International Airport (PHL) was congested about 16% of the time, and that more than 27% of its departures took off when the airport was in a congested state. The resultant taxi-out times, and therefore fuel burn, of these flights in congestion are nearly double their unimpeded values [1].

Observations such as the above have motivated the development of airport surface congestion management strategies. These algorithms range from the aggregate approaches demonstrated at Boston (BOS) and Dallas Fort Worth (DFW), to the aircraft-specific approach tested at New York John F. Kennedy airport (JFK). Each of these approaches has been designed and tailored for one particular airport. However, there is a desire to successfully extend each of these approaches to other airports. System identification, namely, the detailed characterization of the airport operating environment and performance, is a critical step in this process, and is the focus of this paper.

II. RELATED EFFORTS

There have been several efforts in the United States and Europe to develop and implement surface congestion management strategies, especially in the context of Airport Collaborative Decision Making (A-CDM). Examples include the field-testing of the Pushback Rate Control strategy at BOS [2,3], the Tower Flight Data Manager (TFDM) demonstration at Dallas Fort Worth (DFW) airport [4], the field evaluation of the Collaborative Departure Queue Management concept at Memphis (MEM) airport [5], the Surface Congestion Management Program at New York JFK airport [6,7], the human-in-the-loop simulations of the Spot and Runway Departure Advisor (SARDA) concept at DFW [8], the trials of the Departure Manager (DMAN) concept [9] in Athens International airport (ATH) [10], and the Airport Collaborative Decision Making (A-CDM) implementations at London Heathrow (LHR), Frankfurt (FRA), Amsterdam (AMS), Helsinki (HEL) and Paris Charles de Gaulle (CDG) Airports [11]. There has also been increased interest from major airports in Asia (e.g., Bengaluru International Airport (BLR) in India, as well as China and Singapore [12]) in A-CDM. The above surface management approaches can be broadly classified as aggregate approaches that are implemented by the airport tower [2,3,4], airline-specific allocation approaches [5], and aircraft-specific metering approaches [6-11].

Prior surface congestion management efforts have differed both in their implementation details and in the underlying algorithms. The specifics of implementation vary significantly depending on the operating procedures at the airport, the level of automation in the decision support tools, etc. However, it may be possible to deploy the same general congestion management algorithm at multiple airports with some tuning of parameters to suit the particular airport. In order to maximize the return on investment, it is desirable to develop techniques that will aid in the airport-specific adaptation of different surface congestion management algorithms. Such techniques will also enable the evaluation and comparison of different algorithmic approaches.

It is important to note that airport-specific adaptation is essential for the practical implementation of any of the approaches mentioned in the previous section. Each of those

algorithms requires the prediction of airport capacity and throughput, taxi times, and other performance characteristics over a range of operating conditions. This paper addresses this critical need by proposing techniques for the characterization of different airport environments, with the ultimate objective of enabling widespread deployment of surface congestion management algorithms. The proposed techniques combine qualitative observations through site surveys and data visualization, with a quantitative analysis of operational data. The techniques characterize airport operations by identifying aspects that are common across multiple airports, as well as those that differ within the same airport, depending on operating conditions (e.g., runway configuration, weather conditions, demand levels, etc.) Finally, the proposed characterization approaches also help in estimating performance metrics that can be used in evaluating the benefits of different congestion mitigation schemes.

III. DESIGN OF CONGESTION MANAGEMENT APPROACHES

The overall process of designing a congestion management approach is 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. Assuming candidate airports have been selected, each of the subsequent elements are discussed in more detail in the following sections, with the Airport Characterization piece receiving most attention given the focus of this paper.
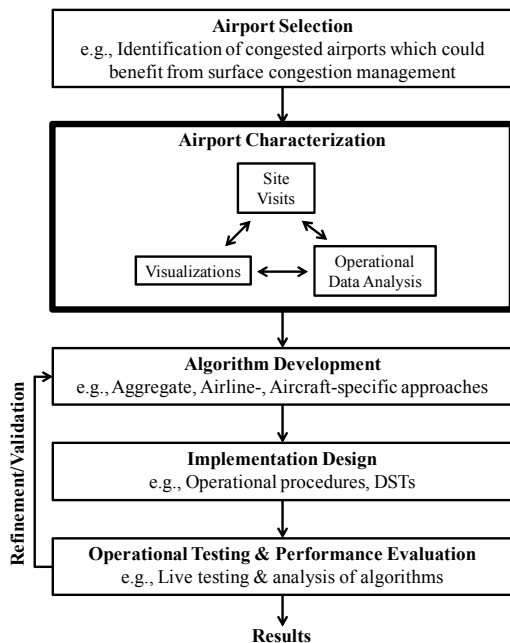

Figure 1. Overall design process for a congestion management approach.

IV. AIRPORT CHARACTERIZATION METHODOLOGY

The three key airport characterization elements described here are site visits, visualizations and operational data analysis, which work together iteratively to help researchers develop an understanding of the characteristics of a given airport. Three airports, Boston (BOS), LaGuardia (LGA) and Philadelphia (PHL), are used as case studies in the discussions in this paper. For reference, the layouts of the three airports are shown in Figure 2.
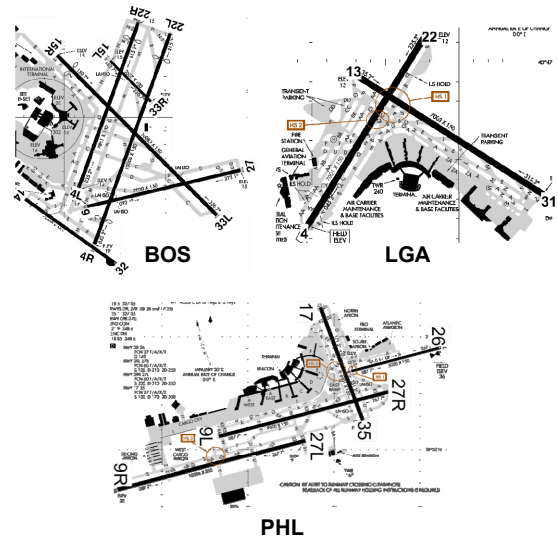

Figure 2. Airport diagrams of BOS, LGA, and PHL airports, showing number and orientations of runways.

A. Site Visits

Site visits allow researchers to get first-hand knowledge of the specific characteristics of individual airports which can be invaluable in determining their feasibility for surface congestion management. Primary benefits of visits include:

- Better understanding of the physical layout of the airport, its equipment levels and operational characteristics such as carrier and fleet mix.
- First-hand observations of operations to better understand standard procedures and current challenges at the airports.
- Ability to gather expert opinions from air traffic control professionals at the airport to help understand operations, get answers to key questions relevant to surface congestion management and to identify potential opportunities for mitigation.

An example of the assessment of physical layout of the airport, ATC tower and equipage levels from a site visit to LGA is provided in Figure 3. The management of pushback processes is of critical importance to surface congestion management techniques, and this figure highlights (in red boxes) the location of airline ramp towers at the airport from which the pushback control occurs. The layout of the FAA air traffic control tower is also illustrated, showing the locations of the personnel responsible for the various tasks in the tower,

what equipment they have available to them, and the nominal flow of the flight progress strips through the tower system (which mirror the physical movement of the aircraft through the airport processes).
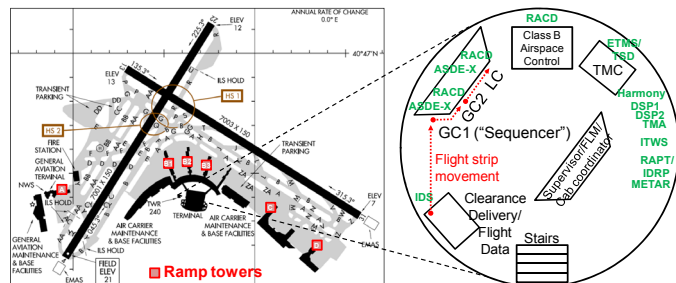

Figure 3.   LGA site visit physical layout and equipage.

An example of the insights that can be gained from site visits to LGA are presented in Figure 4. This shows the nominal arrival (in red) and departure (in green) taxi routes for one of the main configurations at the airport, as well as examples of some of the operational challenges (in magenta dashed line) identified through observations and expert discussions. These highlight where interactions between standard taxi operations (handled in the ATC tower at LGA) and pushback processes (initiated from the ramp towers) may occur and hence help inform the need for, and implementation strategies of, surface congestion management processes at that airport. Similar site visits have been conducted at other locations, including BOS and PHL.
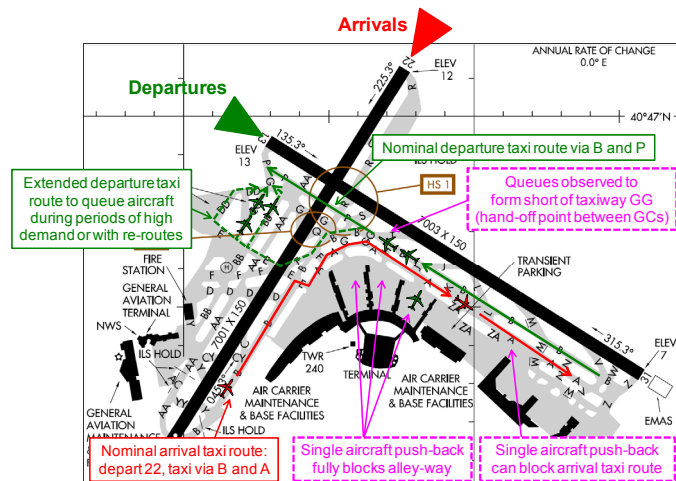

Figure 4.   LGA site visit standard operations and challenge identification.

B.   Qualitative visualization of surface dynamics

Site visits can provide significant insights into operations on a specific day, but a broader understanding of operations across a range of operating conditions (e.g., demand, airport configuration, weather conditions, etc.) can be gained through visualizations of surface traffic. In particular, data from the Airport Surface Detection Equipment, Model X (ASDE-X) surveillance system provides track data for individual flights on the airport surface at a 1 Hz update rate. This data is being archived for a number of airports, which in turn allows detailed traffic visualizations to be created for those sites using appropriate mapping software. This provides a dynamic look into aircraft movements on the surface under different operating conditions to better understand:

- Surface procedures, e.g., in terms of standard taxi routes; runway exit, entry and crossing locations; aircraft holding locations, etc.
- Surface characteristics, e.g., in terms of typical aircraft queuing locations.
- Dynamics of demand as a function of gate, terminal and runway throughout the operating day.
- Dynamics of interactions, e.g., between arrivals and departures on runways, taxiways and ramp areas.

Each of these characteristics is relevant for surface congestion management, either in terms of defining the need for or informing specific implementation approaches. An example of some of these insights is illustrated in the visualization snapshot from LGA operating under configuration 22 | 13 shown in Figure 5.
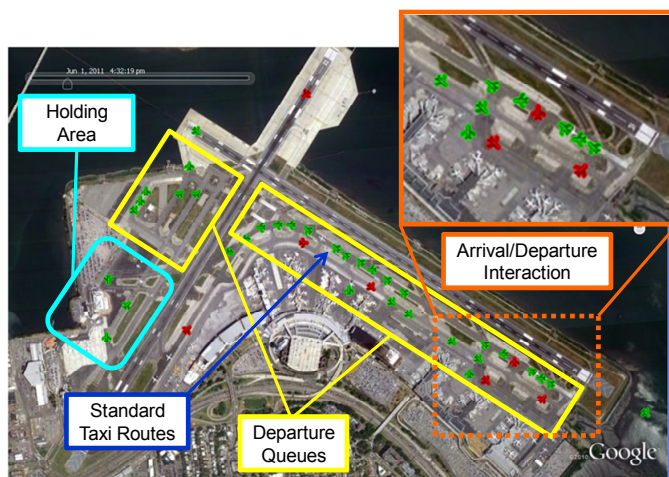

Figure 5.   LGA visualization snapshot.

This is an example of a heavy-demand departure queue for runway 13, with departures again represented by green icons and arrivals by red icons. It highlights several of the airport characteristics identified above. Firstly, the locations of the departing aircraft clearly define the standard taxi routes for this configuration, and confirm the taxi routes identified from the site visit illustrated in Figure 4. Secondly, the location of stationary aircraft identified from the ASDE-X data, shown by the two yellow boxes, allows queue boxes to be identified. Defining queue boxes is an important element of surface congestion management as the numbers of aircraft in these areas are essential control variables. In this period of high demand, the queue extends all the way to the opposite end of the departure runway. The need to cross the arrival runway adds to the congestion and is a dynamic which would also need to be taken into account in any congestion management strategy.

3

Departing aircraft with ground holds add more congestion and are handled differently at every airport. The light blue box is a waiting area for such aircraft. Because LGA does not have much free surface space, the location changes with each configuration.

Lastly, the interaction between arrivals and departures can be seen occurring in the inset. In this example, there are gate conflicts for some arriving aircraft. Since they cannot immediately proceed to their gate, they must instead pass by their gate and wait in the area highlighted in orange. Once the gate is free, an aircraft joins the departure queue and slowly moves back to its gate.

The visualizations for PHL were used to study similar characteristics. One unique finding, specifically in configuration 9R | 9L, is that departure queues during periods of high demand often extended far enough back on the main taxiway to block several of the ramp areas. The visualizations may also lead to more questions. In the process of defining queue boxes it was noted that departure runway 27L has five feeder taxiways, but the visualizations did not reveal any specific method for determining which taxiway the next departing aircraft would come from. Such questions can be noted and then asked during site visits for further clarification.

C. Operational Data Analysis

There are various sources of operational data which can be used to perform analyses of relevance to surface congestion management, including ASDE-X archives and the Aviation System Performance Metrics (ASPM) database. While ASDE-X was originally designed to enhance safety, its high accuracy/update rate allows the surveillance data to be of high value to the airport characterization task. The availability of ASDE-X at the 35 OEP airports in the US makes it particularly attractive for this effort. However, much of the analyses presented below can be carried out with data from the ASPM database, which provides the Out, Off, On and In (OOOI) times of flights in the US National Airspace System. More detailed analyses, such as the measurement of departure queues and runway utilization, require the high-fidelity of ASDE-X data. The following sections present some of the key analyses which can be conducted to inform surface congestion management approaches.

1) Diversity of runway configuration use

In order to adapt a congestion management strategy to a particular airport, it is first important to identify the runway configurations that are most commonly used as this impacts which configurations are most important from a surface congestion management perspective. The frequencies of runway configuration usage at BOS, LGA and PHL in the summer months of 2011 derived from ASDE-X data are presented in Figure 6.

We see that while BOS is predominantly in one of two configurations (4R, 4L | 9, 4R and 22L, 27 | 22R, 22L), while LGA is significantly more heterogeneous in its configuration usage. By contrast, PHL spends nearly 77% of the time in the some variant of the 27R | 27L configuration, since there are only occasional operations on runways 35/17 and 26. From an

adaptation perspective, algorithms at PHL and BOS will need to be tuned for two main configurations in order to be valid over 95% of the time, where as those at LGA will need to adapt to five different configurations in order to be valid to the same extent.
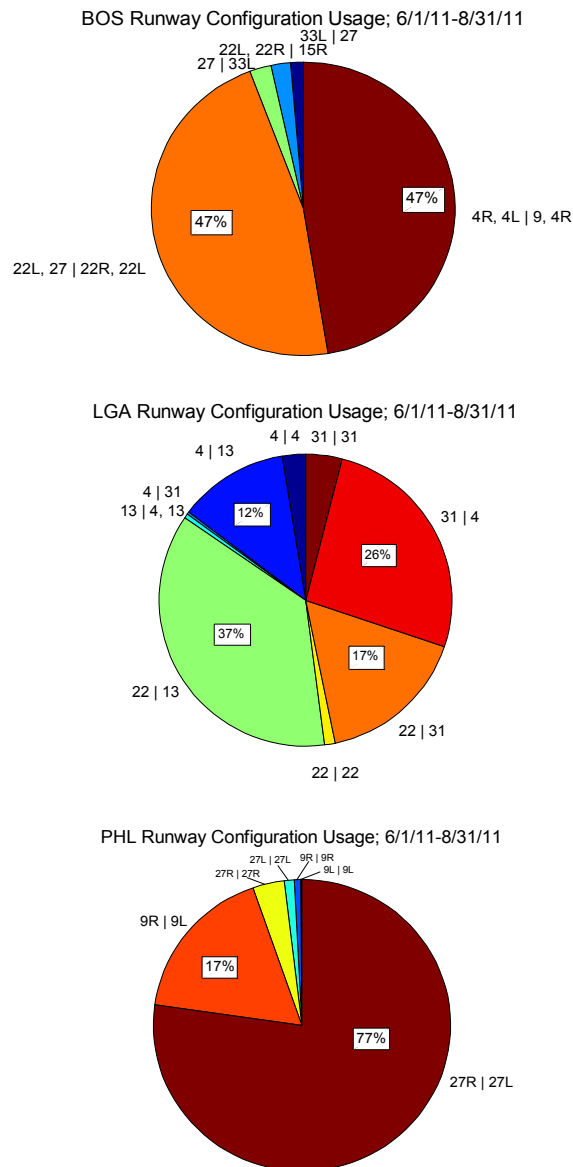

Figure 6. Runway configuration usage at BOS, LGA and PHL.

2) Airline operations mix

As part of airport characterization, it is necessary to study the mix of airline operations at a given airport in order to choose an appropriate congestion management strategy. This data analysis defines who the main airline stakeholders are for each airport. Any sort of air traffic operations change will require involving the relevant airlines in the process and adjusting the algorithm to any specifications or requirements from the airlines. The airline operations mix at BOS, LGA and PHL are shown in Figure 7. In these case studies, we see PHL is a hub for USAirways, with 68% of airport operations

occurring by this dominant carrier. Conversely, BOS and LGA operations are divided among more carriers.

The airline mix at a given airport will influence the type of congestion management strategy chosen. An airport which is dominated by one airline, as in the case of an airline hub, could more easily utilize an airline-specific congestion management strategy, whereas an airport where the operations are more equally split amongst several airlines might be more suitable for an aggregate solution, not tailored to any one specific airline.
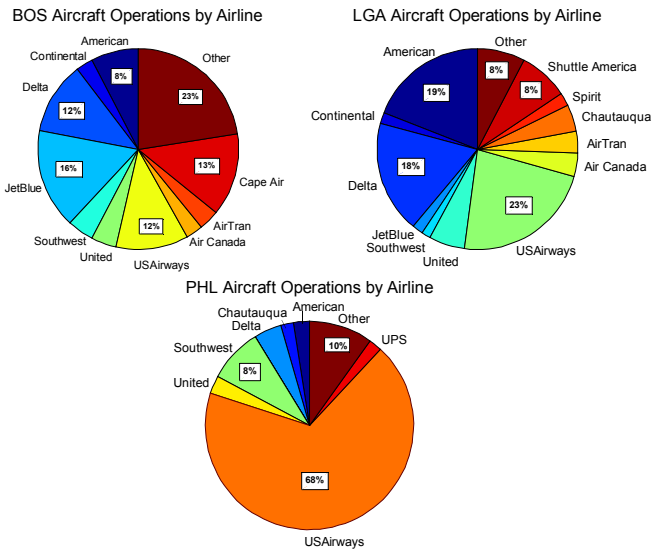

Figure 7. Airline mix at BOS, LGA and PHL.

3) Traffic demand/taxi time characteristics

Because the operations for each configuration can vary significantly, the airport surface dynamics must be analyzed for each configuration separately in order to create the most effective congestion management strategy. Once the dominant configurations have been established for a given airport, the surface operation can be characterized for the dominant configurations. The three figures below show the dynamics of aircraft demand, taxiing, and queuing for the dominant configuration at BOS, LGA and PHL respectively. Studying the magnitude and scope of such metrics can inform both the need for surface congestion management and the potential benefits derived.

The average values for active departures, queue sizes and taxi times over Summer 2011 are shown in Figure 8. The active departures are the number of departing aircraft which have pushed back from the gate but have not yet taken off. We see there is a similarity in the shape of the curves for active departures, queue size and taxi time within each airport. As expected, an increase in departure demand correspondingly increases taxi times and queue sizes.

One characteristic worth noting is the difference in the magnitude and dynamics of demand between each of the airports. We see BOS exhibits relatively low demand throughout the day with congestion peaking at the evening push. LGA exhibits congestion throughout the day, with increased levels seen in the morning and evening. By contrast, PHL demand shows a peak roughly every two hours. This is attributed to PHL being a hub airport for USAirways, exhibiting a banking effect as typical for airline hub connections.
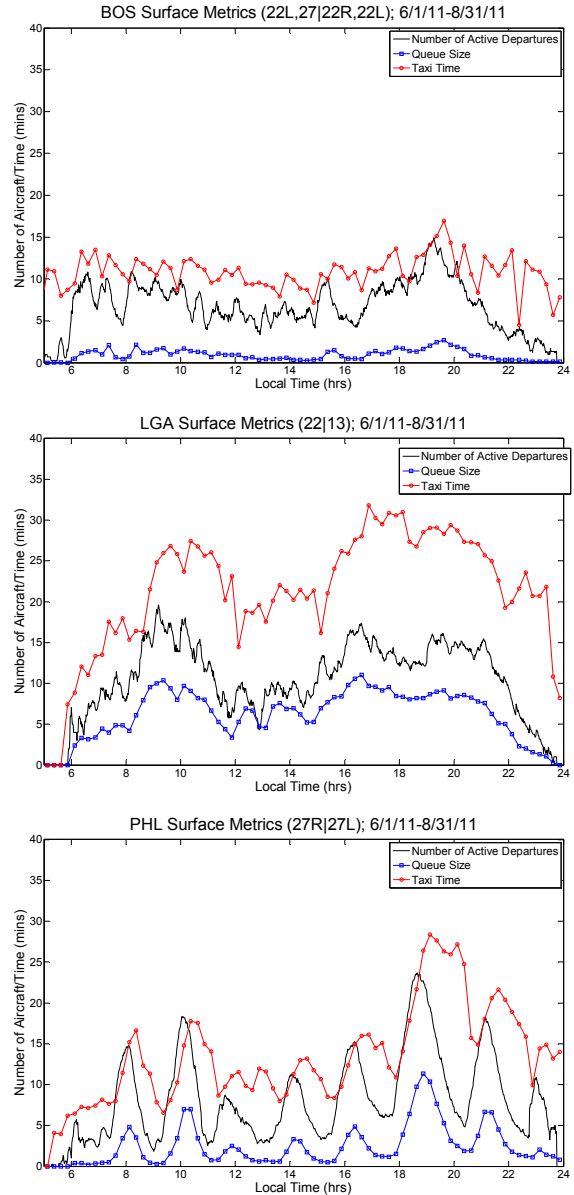

Figure 8. Traffic demand & taxi time characteristics for BOS, LGA & PHL.

Looking at these plots can help inform the need for surface congestion management strategies. At BOS, the periods of congestion are fairly sporadic, with a typical evening time period of congestion. An airport like this may only need to utilize congestion management techniques during a certain time period of the day. LGA, by contrast, has fairly constant congestion throughout the day. Such an airport could benefit from consistent use of congestion management techniques. At PHL, demand rises and falls quickly so a solution might only

be used intermittently here. These differences could be applied more generally as classes of airports for deploying similar congestion management algorithms at multiple airports.

Both the queue sizes and taxi time measurements can be used for predicting potential benefits of various surface congestion management techniques. These are also key metrics to be used in the next phase of algorithm development for tailoring a given strategy to an individual airport.

4) Throughput saturation curves

ASDE-X data is also used to model the throughput characteristics of a given airport. A useful characterization of the throughput performance is to consider the plot of the takeoff rate (or departure rate) as a function of the number of active departures [2,13]. This plot for the 22 | 13 runway configuration at LGA is show in Figure 9. The throughput plots typically exhibit saturation in the departure rate beyond a certain number of active departures, and are therefore referred to as *throughput saturation curves*. Key parameters associated with these plots include the number of active departures corresponding to saturation, and the associated sustained departure throughput (about 11 aircraft and 0.6 departures/min respectively, in Figure 9). We also note that the standard deviation associated with the sustained departure throughput can be quite high, due to the different sources of uncertainty in airport operations, such as, arrival demand, ATC workload, downstream airspace constraints, fleet mix, etc.
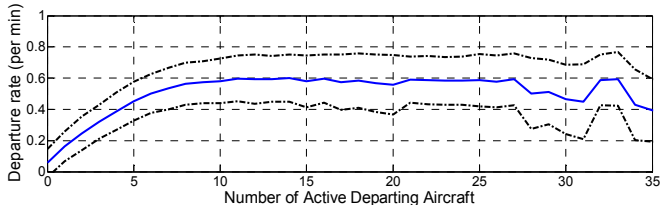

Figure 9. Throughput saturation curve for 22 | 13 runway configuration at LGA. The solid line shows the mean, and the dotted lines show one standard deviation from the mean values.

The throughput saturation curves can differ significantly between two configurations at the same airport, as is illustrated in Figures 10 and 11 for the 9R | 9L and 27R | 27L configurations at PHL. A comparison of the two figures shows that although the same two physical runways are being used in each configuration with the directions reversed, the runway configuration that uses runway 27R for arrivals and 27L for departures achieves a significantly higher departure throughput (mean of 48 departures/hour) that the one that uses 9R for arrivals and 9L for departures (mean of 42 departures/hour). We also note that the 27R | 27L runway configuration saturates at a lower value of active departures than 9R | 9L. One possible explanation for both the above observations is that the use of the outer parallel runway for arrivals increases the number of active crossings of the departure runway, thereby decreasing the departure throughput. The significant differences in the throughput saturation curves implies that the algorithms described in Section V will need to be tuned to different parameters depending on the runway configuration.
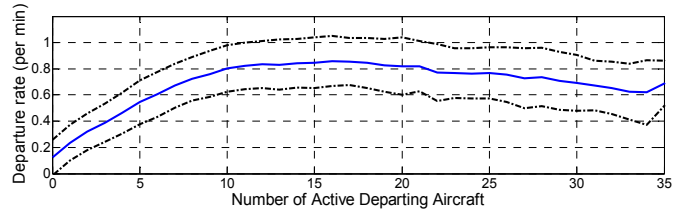

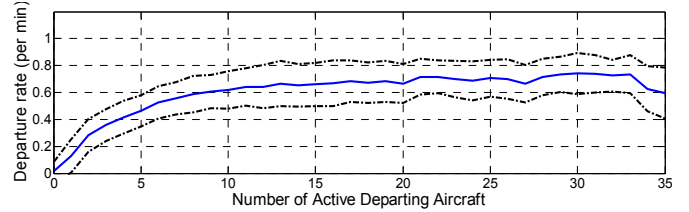Figure 10. Throughput saturation curve for 27R | 27L configuration at PHL.

Figure 11. Throughput saturation curve for 9R | 9L configuration at PHL.

Configurations across different airports that look similar at first glance can in reality have significantly different throughput saturation curves, due to differences in overall airport layout, runway lengths, procedures, etc. Figure 12 shows the throughput saturation curve for the 22L, 27 | 22R, 22L configuration at BOS, which is similar to the 9R | 9L configuration at PHL in the sense that the outer runway is used for arrivals and the inner one is used for departures. There are occasional departures on 22L as well, but 22R and 22L are treated as a single runway in operational practice.
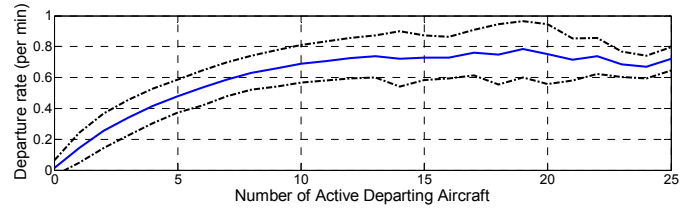

Figure 12. Throughput saturation curve for 22L, 27 | 22R, 22L configuration at BOS.

Despite the similarity between the two sets of operations, we note that the throughput saturation curves (Figures 11 and 12) are considerably different, possibly because of the differences in their operational procedures, airport layout, etc. Throughput saturation curves also illustrate the differences between theoretical and empirical capacity estimates, which can differ due to factors such as airport layout, procedures and human factors [14,15].

Finally, the throughput saturation curves also demonstrate a key benefit of limiting the number of active departures, namely, the deterioration in departure throughput at very high levels of departure traffic. This phenomenon is clearly seen in Figure 10, where the departure throughput initially increases, but then starts to decrease as the number of active departure exceeds 20 aircraft. It is conjectured that very high levels of traffic increase complexity and controller workload, and result in a deterioration of performance.

V. ALGORITHM DEVELOPMENT

The underlying logic in most of the airport management efforts mentioned in Section I relies on determining some indicator of congestion on the surface, and then trying to control pushbacks to maintain that indicator below some value. The level of traffic on the surface and the length of the departure queue are the quantities most often used as an indicator of surface congestion. For instance, the Pushback Rate Control protocol tested in 2010 controlled pushbacks in order to maintain the number of active departures on the surface below a predetermined value, denoted N_{Ctrl} [2]. Other approaches, such as A-CDM, control pushbacks during periods in which the demand exceeds the declared capacity [11]. The algorithms therefore require (1) the prediction of the indicator variable, and (2) determination of the threshold value beyond which pushbacks will be controlled. The threshold value is sometimes determined manually by an expert on the field [5-7], and sometimes from the throughput saturation plots described in the previous section [2,4]. By contrast, the variant of Pushback Rate Control tested in 2011 used operational data to determine the optimal pushback rate as a function of the predicted length of the departure runway queue at the next time period, and used the current observations of the number of taxiing departures and the length of the departure queue to predict the departure runway queue at the next time period [3]. A common feature in nearly all the approaches is the need to accurately predict airport throughput in some time period, which can be done, for instance, using the throughput saturation plots of Section IV, and the observed number of active departures.

Table I presents the values of sustained departure throughput and the number of active departures at which the throughput saturates, for different runway configurations at BOS, LGA and PHL. These values can be determined through an analysis of the corresponding throughput saturation plots.

TABLE I. SUSTAINED DEPARTURE THROUGHPUT AND CORRESPONDING NUMBER OF ACTIVE DEPARTURES, FOR DIFFERENT RUNWAY CONFIGURATIONS AT BOS, LGA AND PHL.

		Saturation point (# active dep.)	Sustained dep. throughput (ac/hr)
BOS	4R, 4L | 9, 4R	17	48
	22L, 27 | 22R, 22L	13	45
LGA	22 | 13	11	36
	31 | 4	15	40
	22 | 31	18	42
	4 | 13	15	36
PHL	27R | 27 L	12	48
	9R | 9L	20	40

After the initial development of a congestion management algorithm, the next step is to assess its potential impacts through simulations. The data analysis presented in Section IV can be used to develop "what-if" simulations, in order to estimate the impacts of the algorithms being considered for implementation [2,16]. In addition to assessing the potential pool of benefits achievable with a particular algorithm (for example, a particular choice of N_{Ctrl} [2]), simulations can also help evaluate potential challenges (such as an increase in gate-use conflicts due to departure metering) prior to implementation.

For example, prior work showed that if the taxi-out times of all flights in congestion periods (that is, they pushed back when the number of departures on the surface exceeded the saturation point) was decreased to their expected value when the airport was at the saturation point, there would be significant reductions in taxi-out times, fuel burn, and emissions [1]. PHL, under such a policy, would benefit from an approximately 13.5% decrease in these quantities, resulting in savings of 2.9 million gallons of jet fuel, nearly 22,000 kg of HC, more than 212,000 kg of CO and more than 37,000 kg of NOx per year. BOS, although less congested than PHL, would benefit from a 6.5% decrease in these impacts, or the estimated savings of 900,000 gallons of jet fuel per year, along with 6,000 kg of HC, 64,000 kg of CO and 11,000 kg of NOx savings annually [1]. Simulations show that LGA, when similarly controlled to its saturation value, would save an estimated 242,000 gallons of fuel in the 31 | 4 configuration alone, even after gate-conflicts with arriving aircraft were resolved. Such simulation-based estimates of the benefits potential of the final refined algorithm can be then be compared to the benefits actually achieved in the operational setting, as described next.

VI. IMPLEMENTATION DESIGN

Once a candidate surface congestion management algorithm has been developed, it is necessary to design implementation strategies which allow it to be tested in an operational setting. The factors which need to be considered when developing an implementation design for any particular airport include:

- Airport/ATC tower operating characteristics.
- Algorithm information input requirements.
- Algorithm execution platform.
- Algorithm output format.
- Algorithm execution procedures.

1) Airport/ATC operating characteristics
This includes factors such as whether the pushbacks are ramp tower or FAA tower-controlled, equipage levels in the various towers, and the layout of the ATC tower. The primary pushback management location is critical because this impacts to whom the congestion management recommendations should be presented. The site visits previously described allow many of these factors to be established early in the assessment of the suitability of an airport for surface congestion management such that any constraints or opportunities afforded at a site can be accounted for.

2) Algorithm information input requirements
Different information input requirements may exist depending on the specifics of the algorithmic implementation. Estimates of capacity and demand at suitable time horizons into the future are of critical importance to all surface

congestion management implementations, as they drive the expected congestion levels. Capacity estimates are largely driven by availability of reliable forecasts of airport configuration and weather conditions (e.g., visibility, ceiling, etc.) which impact airport arrival and departure rates. In terms of demand forecasts, aggregate pushback rate implementation schemes require availability of overall demand estimates over a suitable time horizon, while airline-specific allocation approaches may require the forecast demand to be disaggregated by carrier. Aircraft-specific approaches may additionally require the aggregate demand to be broken down into which exact flights are forecast to be pushing-back in specific time bins into the future. The availability (or otherwise) of these data will play a large role in the type of congestion management algorithm suitable for a given airport.

3) Algorithm execution platform

Many of the field trials of surface congestion management conducted to date have produced their own platform for algorithm execution. For example, Phase 1 of the BOS field trials [2] utilized a paper-based procedure relying on researchers in the ATC tower to gather the appropriate input data (e.g., from visual observation of operations and through other tower computer systems) together with pre-computed throughput saturation curves to manually determine the appropriate pushback rate in 15 minute intervals. Phase 2 of the trials at BOS evolved to use tablet computer devices into which the key input data was entered, and the recommended aggregate pushback rates were then calculated automatically (using the same curves as the manual process). By contrast, the implementations at JFK, MEM and DFW have used dedicated decision support tools (DSTs) in the tower into which key input data is either automatically populated (if suitable electronic feeds are available) or manually entered by controllers or airline personnel.

4) Algorithm output format

The algorithm execution platform generally corresponds to a form of output to manage surface congestion. For example, the initial paper-based procedures at BOS resulted in a congestion management recommendation in terms of a suggested pushback rate communicated by color-coded cards to the gate controller in the ATC tower. This information was subsequently displayed on the tablet device when that was the algorithm execution platform in later trials, while dedicated DST displays were employed in the JFK, MEM and DFW implementations.

5) Algorithm execution procedures

All of the elements above are brought together into specific procedures to execute the chosen algorithm. For example, this would cover the roles and responsibilities for different researchers and/or ramp/airline/ATC tower personnel when conducting surface congestion management with the system. An example of the resulting implementation design used in the Phase 2 BOS trials is shown in Figure 13.
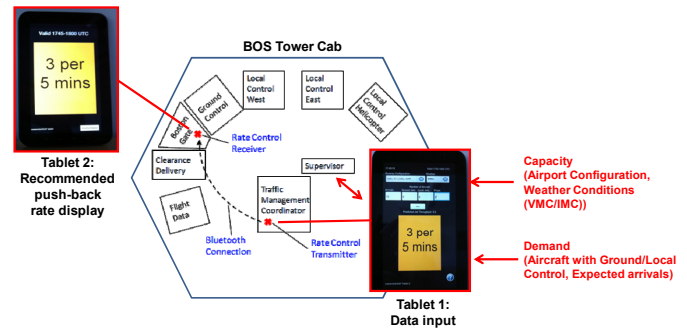

Figure 13. BOS field trial example implementation.

BOS is an FAA tower-controlled pushback airport, where the pushback rate recommendation is made to a gate controller. In this implementation, researchers entered the required algorithmic inputs into a tablet device (shown on the right), which in turn used the appropriate throughput saturation curve (pre-computed from ASDE-X data analysis) to determine the recommended push rate. The researchers in this implementation verified the recommendation with the tower supervisor, after which it was transmitted to a second tablet device located at the gate control position for execution.

It is seen that the aggregate control approach adopted in the BOS implementation could be accomplished entirely from within the FAA tower (albeit with the knowledge of other airport stakeholders). By contrast, schemes developed for JFK and MEM required active real-time participation of many other stakeholders, in particular constantly updated demand information from airlines at those airports.

VII. OPERATIONAL TESTING AND PERFORMANCE EVALUATION

1) Operational testing

Once surface congestion management algorithms have been developed and implementation strategies designed, it is possible to test their performance in an operational setting. The scope of testing opportunities has varied for the different airport activities to date: the DFW operational testing totaled two weeks, the BOS trials conducted testing over two summer periods (35 total test days), while the JFK activities have been tested operationally for over a year. Extensive operational testing is a critical element to ensure validity and robustness to a wide range of operational conditions for any surface congestion management scheme and/or as a basis for algorithm refinement (shown by the feedback in Figure 1).

2) Performance evaluation: Benefits/cost assessment

Operational testing also generates large amounts of data which enable benefits/costs to be estimated. Benefits are typically assessed by comparing appropriate performance measures before and after surface congestion management implementation, with other relevant operational factors being as equal as possible. Given the overall objectives of these activities, assessment of congestion metrics before and after implementation are of particular interest, particularly in terms of taxi times, fuel burn and emissions production. For

example, the two field trials at BOS resulted in a 56-67 kg decrease in fuel burn per gate-held flight (depending on gate APU use), resulting in a total saving of 6,900-8,200 gallons of fuel and 68-80 metric tonnes of CO_2 emissions reduction over the trial period. Decision-makers are often interested in the "monetizable" impact of the deployment of a given system, and in the case of surface congestion management, impacts can be converted into passenger value of time, aircraft direct operating cost and fuel burn costs using standard FAA techniques [17]. There are also newly emerging multipliers to estimate the cost impacts of environmental impacts (e.g., from reduced engine emissions on the ground) as well [18].

3) Performance evaluation: Airport performance

It is also important to assess whether there are any adverse consequences of the approaches, such as loss in departure capacity (discussed here) as well as impacts to other airport stakeholders (airlines, ground crews, etc.). There are several operational constraints specific for airports, which manifest themselves in the throughput characteristics. It is important to account for these effects when designing a departure management strategy. Inter-departure separations which drive throughput depend on several factors such as the weight classes of the leading and trailing aircraft, and their departure fixes. Air traffic controllers generally have formal or informal target inter-departure separations, and these can be compared with actual operations to identify sources of inefficiency. Operational separations can vary both above and below the target values. For example, conservative departure separations can cause the separations to be larger than the targets. On the other hand, certain techniques such as the use of two crossing runways for departures, can allow departures to take place with lesser separations. Figure 14 shows a graphical representation of departure separations at BOS when departures are on the two parallel runways 22R and 22L. Since both runways use the same departure fix, they operate as a single runway. The height of each spike denotes the weight class of the leading aircraft, while the width denotes the target departure separation. It can be seen that the first departure requires a separation of 65 seconds after it. However, there is an additional 15-second gap before the next departure actually occurs. On the other hand, the next two departures happen precisely at the completion of the target departure separations. It should be noted that the width of each spike also depends on the weight class of the trailing aircraft, the runway configuration as well as airport-specific considerations. However, the procedure is the same for identifying opportunities for improving departure efficiency. When consistently tracked over several months, departure spacing can be used to measure the average operational performance of the airport [19].
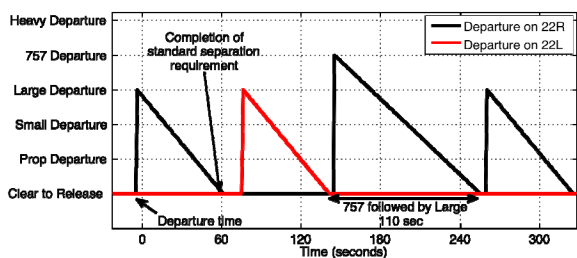

Figure 14. Successive departures from runways 22R and 22L at BOS.

In addition to assessing nominal airport performance, departure spacing plots can be used to evaluate the performance of surface management strategies, in particular, to identify periods when the runway was starved (a highly undesirable state if demand exists). By investigating time periods in which departure separations were larger than usual, one can determine whether the increased gaps were due to a lack of aircraft in the departure runway queue, and whether departure metering was responsible for the starvation of the runway queue [16].

Another metric of airport performance is runway utilization. If runway utilization is maintained during departure metering, it is reasonable to expect that gate-hold times will translate into taxi-out time reduction. Runway utilization is estimated by determining (using ASDE-X data) what percentage of each 15-min interval corresponded to a departure on takeoff roll, aircraft crossing the runway, arrivals (that requested landing on the departure runway) on final approach, departures holding for takeoff clearance, etc. Figure 15 shows the runway utilization on an evening during the Pushback Rate Control trials at BOS in 2010 [16]. It is seen that between 1745 and 2000 hours, when gate-holds were experienced, the runway utilization of the primary departure runway (33L) was maintained near 100%, with a persistent departure queue as well.
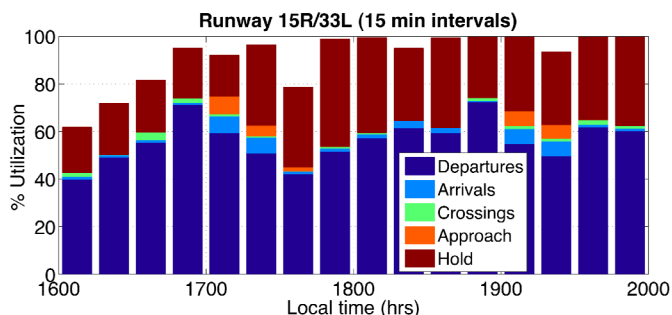

Figure 15. Runway utilization of the primary departure runway (33L) during field-tests of the Pushback Rate Control strategy in 2010 [16].

VIII. SUMMARY

This paper laid the groundwork necessary for the successful implementation of surface congestion management strategies, namely the characterization of the airport operating environment. The proposed approach to airport characterization combined site surveys and the qualitative visualization of surface surveillance data with the quantitative analysis of operational data from various sources. The analysis techniques proposed incorporate factors such as airport and tower cab layouts, locations of surface queues, runway configurations, throughput saturation and taxi times under each configuration, traffic demand patterns, air carrier mix, etc. In addition to tuning the parameters of the congestion control algorithms, the paper also proposed metrics for evaluating the impact of field tests and implementation. It is believed that the establishment of approaches for airport-specific adaptations will be essential for the widespread deployment of surface congestion management strategies, and ultimately the mitigation of the fuel burn, emissions and noise impacts of airport surface operations.

REFERENCES

[1] I. Simaiakis and H. Balakrishnan, "Impact of congestion on taxi times, fuel burn, and emissions at major airports," Transp. Res. Rec. No. 2184, Transportation Research Board of the National Academies, December 2010.

[2] I. Simaiakis, H. Khadilkar, H. Balakrishnan, T. G. Reynolds, R. J. Hansman, B. Reilly and S. Urlass, "Demonstration of Reduced Airport Congestion through Pushback Rate Control," Proceedings of the 9th USA/Europe ATM R&D Seminar, June 2011.

[3] I. Simaiakis and H. Balakrishnan, "Dynamic Control of Airport Departures: Algorithm Development and Field Evaluation," American Control Conference, July 2012.

[4] H. Reynolds, M. Ishutkina, D. Johnson, R. Jordan, M. Kuffner, K. Lokhande, T. Reynolds and S. Yenson, "Dallas/Fort Worth Field Demonstration #2 (DFW-2) Final Report for Tower Flight Data Manager (TFDM)," Project Report ATC-393, MIT Lincoln Laboratory, Lexington, MA, 2012.

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

[6] A. Nakahara, T. Reynolds, T. White, and R. Dunsky, "Analysis of a surface congestion management technique at New York JFK airport," AIAA Aviation Technology, Integration and Operations (ATIO) Conference, Virginia Beach, VA, September 2011.

[7] S. Stroiney, H. Khadilkar and H. Balakrishnan, "Ground Management Program at JFK Airport: Implementation and Benefits," Transportation Research Board Meeting, January 2012.

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

[9] D. Böhme, "Tactical departure management with the Eurocontrol/DLR DMAN," in 6th USA/Europe Air Traffic Management Research and Development Seminar, Baltimore, MD, 2005.

[10] M. Schaper, G. Tsoukala, R. Stavrati and N. Papadopoulos, "Departure flow control through takeoff sequence optimisation: Setup and results of trials at Athens airport," in Digital Avionics Systems Conference (DASC), 2011.

[11] European Airport Collaborative Decision Making (A-CDM). http://www.euro-cdm.org/.

[12] E. Lagios, "Airport CDM - Lessons learnt & challenges ahead," META-CDM workshop, European Community's Seventh Framework Programme (FP7), London Heathrow Airport, January 2013.

[13] N. Pujet, B. Delcaire and E. Feron, "Input-output modeling and control of the departure process of congested airports," AIAA Guidance, Navigation, and Control Conference and Exhibit, Portland, OR, 1999.

[14] M. S. Alcabin, R. W. Schwab, M. L. Coats, M. E. Berge and L. S. Kang, "Airport capacity and NAS-wide delay benefits assessment of near-term operational concepts," 6th AIAA Aviation Technology, Integration and Operations Conference (ATIO), September 2006.

[15] I. Simaiakis, "Analysis, Modeling and Control of Airport Departure Processes," PhD Thesis, Massachusetts Institute of Technology, 2013.

[16] I. Simaiakis, H. Khadilkar, H. Balakrishnan, T. G. Reynolds and R. J. Hansman, "Demonstration of Reduced Airport Congestion through Pushback Rate Control," submitted.

[17] FAA, "Economic Information for Investment Analysis Data Package", March 2011.

[18] Dorbian, C., P. Wolfe & I. Waitz, "Estimating the Climate and Air Quality Benefits of Aviation Fuel and Emissions Reductions", Atmospheric Environment ,Vol. 45, pp. 2750-2759, 2011.

[19] H. Khadilkar and H. Balakrishnan, "Estimation of Aircraft Taxi-out Fuel Burn using Flight Data Recorder Archives," AIAA Guidance, Navigation and Control Conference, August 2011.

AUTHOR BIOGRAPHIES

Melanie Sandberg is associate staff in the Air Traffic Control Division of MIT Lincoln Laboratory. She obtained her Master of Science in Transportation from the Massachusetts Institute of Technology. Her research interests include operational strategies for improvements in air traffic management.

Tom G. Reynolds is technical staff in the Air Traffic Control Division of MIT Lincoln Laboratory. He obtained his Ph.D. in Aerospace Systems from the Massachusetts Institute of Technology. His research interests span air transportation systems engineering, with particular focus on air traffic control system evolution and strategies for reducing environmental impacts of aviation.

Harshad Khadilkar is a graduate student in the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology. He received his Bachelors degree in Aerospace Engineering from the Indian Institute of Technology, Bombay. His research interests include algorithms for optimizing air traffic operations, and stochastic estimation and control.

Hamsa Balakrishnan is an Associate Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology. She received her PhD in Aeronautics and Astronautics from Stanford University. Her research interests include ATM algorithms, techniques for the collection and processing of air traffic data, and mechanisms for the allocation of airport and airspace resources.