Network Congestion Control of Airport Surface Operations

Harshad Khadilkar[*] and Hamsa Balakrishnan[†]

Massachusetts Institute of Technology, Cambridge, MA 02139, USA

This paper presents a novel approach to managing the aircraft taxi-out process at airports, by posing the problem in a network congestion control framework. A network model of the airport surface with stochastic link travel times is first developed. The model parameters for Boston Logan International Airport are estimated using empirical data from a surface surveillance system. Using the entry times of aircraft into the network as the control variables, a control algorithm that aims to reduce taxi-out times is proposed. This algorithm regulates network congestion while limiting the adverse effect on throughput performance. Performance characteristics of the proposed control algorithm are derived using simplified theoretical analysis, and are shown to agree with simulations of traffic on a network model of Boston Logan International Airport.

Nomenclature

ASDE-X   Airport Surface Detection Equipment, Model X

BOS   Boston Logan International Airport

$\mathcal{C}$ Cost function for control algorithm

FAA   Federal Aviation Administration

$N_{s,l}$ Number of stops on link l, modeled as a geometric random variable with parameter p_{kl}

a Weight in the cost function on pushback delay, as compared to expected taxi time

k Current level of traffic (number of active or post-pushback departures) on the surface

k_{ctrl} Target number of active or post-pushback departures on the surface

k_p Projected traffic level after time t_p

---

[*]PhD Candidate, Department of Aeronautics and Astronautics. 77 Massachusetts Avenue, Room 35-217, Cambridge, MA 02139. AIAA Student Member.

[†]Associate Professor, Department of Aeronautics and Astronautics. 77 Massachusetts Avenue, Room 33-328, Cambridge, MA 02139. AIAA Senior Member.

$p_{i,j}$	Probability of transition from traffic state i to j
t_l	Travel time over link l
t_p	Pushback delay assigned to an aircraft that calls for pushback at the current time
$t_{s,l,i}$	Stationary time corresponding to i^{th} stop on link l, modeled as exponential random variable with rate μ_l
t_t	Expected taxi-out time of an aircraft
$t_{u,l}$	Unimpeded travel time over link l, modeled as an Erlang random variable with order n_l and rate λ_l
Δt_{kl}	Average inter-departure time from link l when surface traffic level is k
$\frac{1}{\mu_l}$	Expected taxi time added per stop on link l
$\frac{X_l}{\mu_l}$	Expected increase in taxi time on link l due to the addition of one more aircraft on the airport surface
σ_l	Maximum throughput value for link l
ζ_l	Arrival rate to link l

I. Introduction

The reduction of taxi-out times at airports has the potential to substantially reduce delays and fuel consumption on the airport surface, and to improve the air quality in surrounding communities. The taxiway and runway systems at an airport determine its maximum possible *departure throughput*, or the number of aircraft departures that it can handle per unit time. Current air traffic control procedures allow aircraft to push from their gates and enter the taxiway system as soon as they are ready. As this *pushback rate* approaches the maximum departure throughput of the airport, runway queues grow longer and surface congestion increases, resulting in increased taxi-out times.

The FAA's Aviation System Performance Metrics (ASPM) database estimates that annually, taxi-out delays (the difference between actual and unimpeded taxi-out times) at major airports in the United States exceed 32 million minutes.[1] This observation suggests that there are large potential benefits to be attained by implementing congestion control strategies at airports. While there have been several studies with this goal,[2–7] they typically either fail to consider the inherent stochasticity of aircraft movement and pilot behavior, or are intrinsically numerical in nature. By contrast, this paper presents a formulation that explicitly accounts for stochastic taxi-out behavior, while providing insight into the system via an analytical treatment of airport performance. The approach proposed in this paper represents a mesoscopic model, which tries to strike a balance between macroscopic queuing models[2–4] and microscopic, aircraft trajectory-based ones.[6,7] To the best of our knowledge, this paper is the first to develop a model of airport operations using actual surface surveillance data.

A. Related work

Improving the efficiency of the departure process at airports has been a problem of interest for several years.[4–7] Since surface operations are highly heterogenous in terms of taxi speeds, routes and pilot behavior, Eulerian models[8,9] of traffic flow are not applicable here. Most traditional solution approaches involve the use of optimal scheduling algorithms[6,7] or queuing theory,[3–5] both of which have certain drawbacks. Optimization using linear or mixed-integer programming assumes that aircraft move at constant velocities at all times and follow time-based taxi instructions exactly, which are not realistic assumptions given the current state of technology at airports. Ground taxi is controlled manually by the pilots, and furthermore, the voice communication channel between pilots and air traffic control is highly congested, making it difficult to add the command and confirmation messages required for detailed control. Data communication channels are still used only for basic functions, and voice confirmation is considered mandatory for safety reasons. Therefore, the control of pushback times is the most viable near-term strategy.

Queuing theory based approaches to the surface traffic management problem either make simplifying assumptions to obtain analytical results, or else require numerical solutions. In addition, a service process needs to be defined at every queuing node; models typically assume that the queuing occurs only at the runway and not on the taxiways.[3] By contrast, posing the problem in a network congestion control framework offers the advantage of being able to accommodate complex link travel time distributions. At the same time, it allows one to address the issues of network stability and performance through analytical approaches. It also has the advantage of not requiring an explicit service process. Prior data-driven modeling approaches have involved aggregate, queuing models, primarily because the only available data were the gate-out, gate-in, taxi-out and taxi-in times.[1] However, the recent deployment of surface surveillance systems at major airports has enabled the development of network models of the form studied here[10].

The airport surface congestion problems considered in this paper are similar to those encountered in wireless networks,[11,12] TCP congestion control,[13,14] and manufacturing systems.[15–17] One key difference is that, while TCP requires decentralized control, centralized control with global information is possible for airport congestion management. Network models proposed for urban transportation problems[18,19] are suited to networks with a large number of nodes and links, while the airport model contains only a few of each. Airport models also require detailed models of link travel times, and greater emphasis is placed on the optimal time of entry into the airport network than on optimal routing. In the context of airport operations, aggregate rate-control approaches that aim to stabilize surface traffic at a specific level have been proposed.[20] These algorithms are easy to implement in practice, but do not realize the full potential fuel savings and performance gains that may be possible though more detailed congestion management strategies. The solution approach proposed in this paper tries to decrease the level of aggregation in order to increase the benefits from congestion control, while only recommending pushback times to aircraft.

II.   Modeling of taxi-out times

A. Surface surveillance data

Airport Surface Detection Equipment, Model-X (ASDE-X) is primarily a safety tool designed to mitigate the risk of runway collisions.[10] It incorporates real-time tracking of aircraft on the surface to detect potential conflicts and monitor conformance.[21] There is potential, however, to use the data generated by it for surface operations analysis and modeling of aircraft behavior. Reported parameters in ASDE-X include each aircraft's position, velocity, altitude and heading. The update rate is 1Hz for each individual flight track. The main dataset used in this study consisted of 45 days of operations from May-June 2011 at Boston Logan International Airport (BOS), split into a training data set (30 days) and a test data set (15 days). A total of 24,636 departing aircraft were included, and the raw surface tracks were processed using a multimodal unscented Kalman filter developed in prior work.[22] Additional validation of the identified model parameters was conducted using ASDE-X data from February-August 2011.

B. Modeling framework

Figure 1 shows the runways and taxiways at BOS that are represented in the network model. The taxiways form the links of the network, and their major intersections are marked as the nodes. The taxi-out phase for an aircraft is defined to begin when an aircraft leaves the gate, and end when it starts its takeoff roll from the runway threshold. Therefore, the source nodes in the network are the ones adjoining the gates, while the destination/terminal nodes are the runway thresholds. The resulting network abstraction is shown in Figure 2 (a). While the figure shows the union of the networks for all possible airport configurations (allocation of runways to landings and takeoffs), only one configuration is active at a time, and each aircraft has a unique source node and terminal node.

Figure 1. Layout of the surface at BOS

Figure 2 (b) shows the effective network layout when the airport is performing departures from Runway 27. In any specific configuration such as this one, aircraft maintain a flow from the gates to the runway, and generally do not
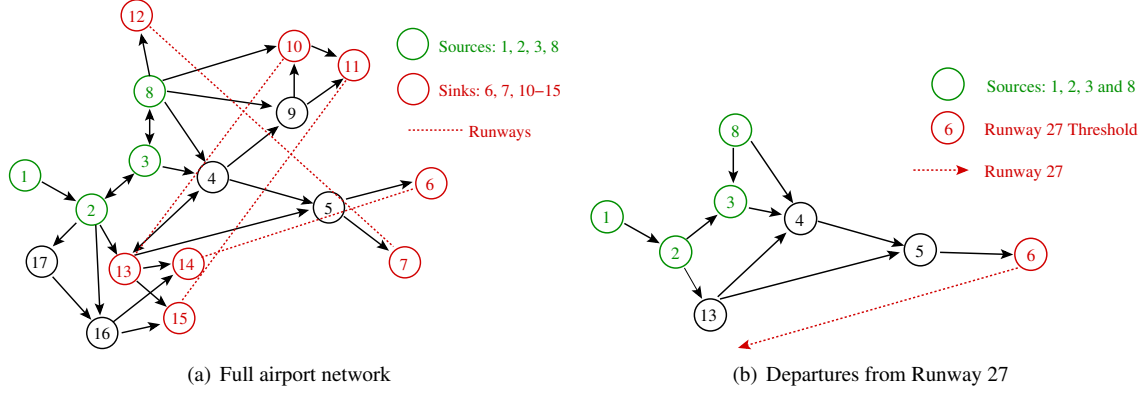
(a) Full airport network   (b) Departures from Runway 27

Figure 2. Abstraction of the airport surface as a graph

taxi in cyclic paths. Consequently, the configuration-specific networks are directed acyclic graphs with random link travel times. Since arrivals taxiing-in at Boston Logan airport tend not to interact with departures taxiing-out, these are not specifically addressed in the proposed model. However, the stochasticity introduced by the occasional interactions between arrivals and departures is captured by the probability distribution of link travel times.

C. Parameter estimation

Aircraft on the surface taxi at fairly constant velocities, occasionally stopping because of other aircraft crossing their path, or when about to cross an active runway. Notionally, the taxi-out process can thus be classified into two modes: *unimpeded taxi*, and *stopped*. Empirical distributions derived from surface surveillance data were used to postulate theoretical models for these two modes. There are three physical processes that govern the taxi time of each aircraft. Firstly, the unimpeded portion of the taxi is a direct result of the length of the link and the velocity at which the aircraft moves. The number of stops that an aircraft makes on each link depends on the congestion on the surface. Finally, the total time that an aircraft spends in the stopped mode also depends on the duration of each stop.

Three sets of theoretical distributions were thus selected to explain the empirical data. The distribution of the number of stops and the duration of each stop showed clear characteristics of geometric and exponential random variables respectively. In the case of the unimpeded taxi time distribution, three candidate theoretical models were considered: Gaussian, log-normal and Erlang. Optimal parameter values for each distribution were calculated by minimizing the Kullback-Leibler (KL) divergence from the true distribution to the model.[23] This resulted in the information projection (p_G^*, p_L^* and p_E^* respectively) of the empirical distribution (q) on each model family. The procedure is illustrated in Figure 3, with the notation $D(p||q)$ denoting the KL divergence from p to q. It was seen that the Erlang family was consistently closest to the empirical distribution for all links in the network, with $D(p_E^*||q) < D(p_L^*||q) < D(p_G^*||q)$. Therefore, the unimpeded taxi times were modeled as Erlang random variables.
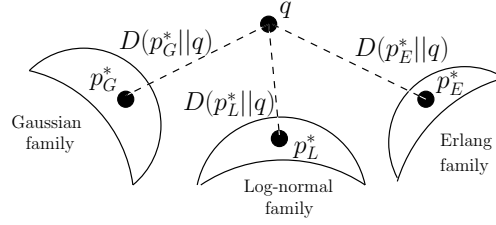
Figure 3. Information projection of the empirical distribution q on three candidate model families

D. Model for link travel times

The travel time over a link l is modeled as,

$$t_l = t_{u,l} + \sum_{i=1}^{N_{s,l}} t_{s,l,i} \tag{1}$$

Here, $t_{u,l} > 0$ is the unimpeded travel time over the link l, an Erlang random variable with order n_l and rate λ_l. $N_{s,l} \in \{0, 1, 2, ...\}$ is the number of stops on the link, modeled as a geometric random variable with parameter $p_{kl} \in [0, 1]$, where k is the current level of traffic on the surface and l is the current link. Finally, $t_{s,l,i} > 0$ is the stationary time corresponding to the i^{th} stop on link l, modeled as an exponential random variable with rate $\mu_l > 0$. The values of $t_{s,l,i}$ are assumed to be independent and identically distributed (*i.i.d.*).

If the number of stops is $N_{s,l} = 0$, then $t_l = t_{u,l}$. Further, each instance of travel time on a link is independent of all other instances, whether on the same link or on other links, conditioned on the level of surface traffic. This observation was derived from empirical data, which showed that surface traffic level as a whole had a greater influence on the taxi time of a given aircraft than aircraft in the immediate vicinity (such as on the same link). Since most departing aircraft on the surface taxi towards a common destination (the runway), there are few instances of interaction between successive aircraft on a given taxi path. The effect of such instances is captured by conditioning on the total surface traffic level. A sample comparison between the empirical and the theoretical probability densities for taxi times on several links marked in Figure 1, is shown in Figure 4. The empirical distribution used is from a set of days independent of those used for training the model. Similar matches were seen between the empirical data and the model for all frequently used links.

The consistency of parameter estimates was tested in order to account for possible seasonal variations in taxi behavior at the airport. An extended data set was used for this purpose, consisting of seven months of ASDE-X data from February to August 2011. The parameter estimation process was carried out for one month of data at a time. Figure 5 shows the results for the unimpeded taxi parameters n_l and λ_l, and the stop time duration μ_l, for a subset of links. The estimated values for each link had fairly small standard deviations, with the exception of link 8→10. Further investigation revealed that the there was a significant deviation for link 8→10 in only one month of data,

while the parameter estimates for the other six months were similar. This anomaly is therefore believed to be due to a lack of sufficient data points in that month. Cumulatively, the above results imply that the network model is a good representation of the actual processes on the airport surface.
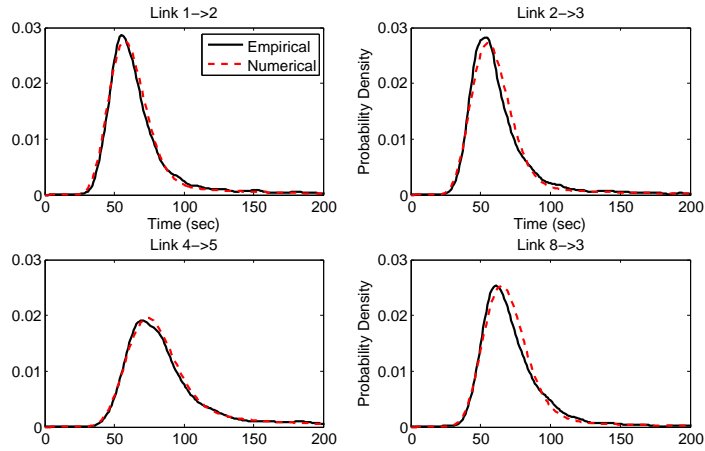

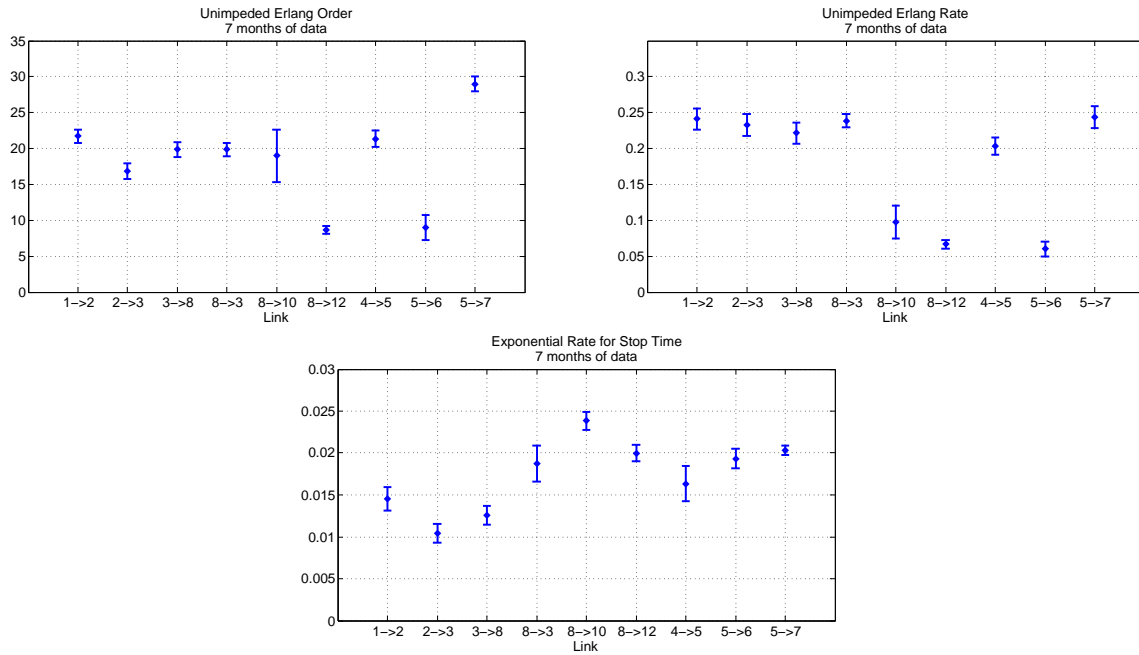**Figure 4. Comparison of probability density of taxi times on some of the major links at BOS**

Figure 5. Means and standard deviations of the parameters for several links at Boston airport

E. Parameter variation with surface traffic

Taxi-out times at airports increase with increased surface congestion, and this needs to be accounted for by the model. The *surface traffic level*, k, is defined as the total number of departures that have pushed back from their gates but have

not taken off yet. Empirical evidence shows that the unimpeded travel time parameters n_l and λ_l, as well as the stop time parameter μ_l, remain invariant with changes in k. The additional taxi-out time due to congestion is accounted for by an increase in the stopping probability on each link, p_{kl}. This fact implicitly accounts for the 'departure queue' that forms at the runway, such as on link $5 \rightarrow 7$ in Figure 1. An increase in the value of p_{kl} will lead to an increase in the number of stops, and each additional stop will add, on average, a penalty of $\frac{1}{\mu_l}$ to the queuing and taxi-out times. The average travel time for each link l can be calculated by taking the expectation of both sides of Equation (1).

$$\mathbb{E}[t_l \mid k] = \frac{n_l}{\lambda_l} + \mathbb{E}[N_{s,l}|k]\frac{1}{\mu_l} = \frac{n_l}{\lambda_l} + \frac{p_{kl}}{1 - p_{kl}}\frac{1}{\mu_l}. \tag{2}$$

Furthermore, assume that each additional aircraft on the surface adds a fixed time penalty to the expectation in Equation (2). Denoting this penalty as a fraction X_l of the expected stop time $\frac{1}{\mu_l}$, it may be observed that for an aircraft that pushes from its gate at a traffic level of $(k+1)$ instead of k,

$$\mathbb{E}[t_l \mid k+1] = \frac{n_l}{\lambda_l} + \frac{p_{kl}}{1 - p_{kl}}\frac{1}{\mu_l} + \frac{X_l}{\mu_l}. \tag{3}$$

Comparing Equation (3) with Equation (2) evaluated at $(k+1)$, the following recursive relationship is seen:

$$\begin{aligned}
\frac{p_{k+1,l}}{1 - p_{k+1,l}} &= \frac{p_{kl}}{1 - p_{kl}} + X_l. \\
\frac{p_{kl}}{1 - p_{kl}} &= \frac{p_{0l}}{1 - p_{0l}} + kX_l \\
\implies p_{kl} &= \frac{p_{0l} + kX_l(1 - p_{0l})}{1 + kX_l(1 - p_{0l})}.
\end{aligned} \tag{4}$$

As shown in Figure 6 for the link $5 \rightarrow 7$, the variation of p_{kl} as derived from Equation (4) is consistent with empirical data, where the probability of stopping was independently measured by looking at the fraction of aircraft that stopped on a link, as a function of the level of traffic on the surface when the aircraft entered the link.
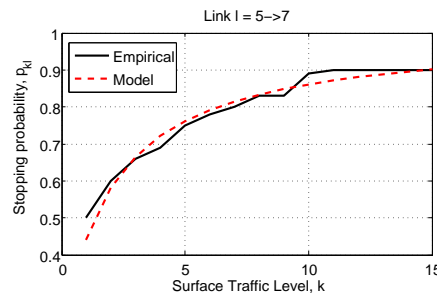

Figure 6. Variation of stopping probability on link $5 \rightarrow 7$ with the level of surface traffic

III. Development of control strategy

A. Derivation of maximum link throughput

Before deriving the maximum throughput of a generic network, it is necessary to find the maximum throughput of a network containing a single link l. Initially, let the link be capable of accommodating an infinite number of aircraft. Furthermore, let it be operating in deterministic steady state, with the taxi time of each aircraft being equal to its expectation value. Aircraft are distributed regularly along its length, with each *departure from* the link occurring after a fixed time interval, and each *arrival to* the link happening at the same instant. If ζ_l is the arrival rate to the link, under this assumption, a new arrival occurs every $\left(\frac{1}{\zeta_l}\right)$ seconds. In steady state, the number of aircraft on the link will be maintained at some value k. From Equation (3), note that the (deterministic) taxi time on the link is given by

$$t_{kl} = \mathbb{E}[t_l|k] = \frac{n_l}{\lambda_l} + \frac{p_{0l}}{(1 - p_{0l})\mu_l} + \frac{kX_l}{\mu_l} = \eta_l + \frac{kX_l}{\mu_l}, \tag{5}$$

where η_l is a constant comprised of the first two terms. During this time interval in which an aircraft travels over the link, the k aircraft ahead of it depart from the link. Therefore, the inter-departure time Δt_{kl} is

$$\Delta t_{kl} = \frac{t_{kl}}{k} = \frac{\eta_l}{k} + \frac{X_l}{\mu_l}. \tag{6}$$

In steady state, k is the result of equating the inter-arrival interval $\left(\frac{1}{\zeta_l}\right)$ to Δt_{kl}. Consequently, the minimum inter-arrival interval that can be sustained by the link is $\left(\frac{1}{\zeta_l}\right)_{\min} = \frac{X_l}{\mu_l}$, achieved as $k \to \infty$. The maximum sustained throughput of link l is the inverse of this value, that is,

$$\sigma_l \triangleq (\zeta_l)_{\max} = \frac{\mu_l}{X_l}. \tag{7}$$

In the stochastic case, the average inter-departure times will be governed by the expectation in Equation (5). Therefore, the result from Equation (7) still holds. Relaxing the assumption of infinite link capacity is also quite simple. Assuming the maximum capacity of the link to be $k_{\max}$, the minimum sustained inter-departure time, as derived from Equation (6), will be $\left(\frac{1}{\zeta_l}\right)_{\min} = \frac{\eta_l}{k_{\max}} + \frac{X_l}{\mu_l}$. The maximum throughput will be the inverse of this quantity. In further analysis in this paper, only infinite-capacity links are considered, with the knowledge that all results may be extended for finite-capacity links. If the input rate to a link l is less than σ_l, the link will be referred to as being *stable*. Once the maximum throughput values for each link are known, it is straightforward to derive the maximum network throughput, using the mincut/maxflow theorem.[24, 25] Note that in general, simply ensuring that each individual link remains stable does not

ensure stability of the entire network.[15, 16] However, it is known that for the special case of directed acyclic graphs, such stability is guaranteed.

B. Definition of control strategy

A control algorithm that aims to maintain a steady level of traffic on the airport surface, is described here. It has been shown[20] that regulating traffic on the airport surface to a well-chosen level results in fuel savings. Therefore, this objective is a logical starting point for any new control strategy. Consider a scenario where an aircraft calls ready to pushback at time $t = 0$. Now consider a First-Come-First-Served (FCFS) algorithm that aims to minimize a weighted linear cost function for the aircraft, composed of *pushback delay* t_p and *expected taxi-out time* t_t (equal to $\mathbb{E}[t_l | k]$). That is, the cost function is $\mathcal{C} = a \cdot t_p + t_t$, with a being a constant weighting factor. Its relationship with the target traffic level is developed towards the end of this section. The control value t_p that minimizes expected cost is given by

$$
\begin{aligned}
t_p &= \arg\min_{t_p \geq 0} \mathbb{E}\left[a\, t_p + t_t\right] = \arg\min_{t_p \geq 0} \left(a\, t_p + \sum_{l=l_1}^{l_r} \left[\frac{n_l}{\lambda_l} + \frac{p_{k_p l}}{1 - p_{k_p l}} \frac{1}{\mu_l}\right]\right) \\
\Rightarrow t_p &= \arg\min_{t_p \geq 0} \left(a\, t_p + \sum_{l=l_1}^{l_r} \left[\eta_l + \frac{k_p X_l}{\mu_l}\right]\right),
\end{aligned}
\tag{8}
$$

where the last two steps follow from Equation (2) and the definition of η_l. Equation (8) assumes that the aircraft's route follows links $l = l_1, l_2, ..., l_r$, and that $p_{k_p l}$ is the stopping probability on link l, when the projected surface traffic level at t_p is k_p. The expected taxi time is thus also a function of t_p. The projected traffic level, k_p, can be calculated based on the procedure outlined below. Since entry into the network is assumed FCFS, the projected traffic level decreases as t_p increases (there can be no additional aircraft entering the network while the current aircraft is waiting). Also, Equation (6) shows that the expected time between departures increases as k_p decreases. Therefore, at some value of k_p, the increase in expected cost due to the first term in Equation (8) is going to outweigh the decrease in cost due to a smaller k_p in the second term. The actual value of this 'target' k_p is controlled by the weight a.

B.1 Simplified model of departure process from the network

To develop a control strategy from an analytical perspective, some simplifying assumptions about the departure process need to be made. The following derivation will be helpful in providing intuition about the departure process. Furthermore, simulations carried out using the full-scale model with no simplifying assumptions (described in the next section) show that the ensuing control strategy is still valid. For simplicity, first consider the single-link network. For moderately large values of surface traffic k, it may be assumed that departures from the link occur as exponential

processes, one for each aircraft. If there are k aircraft on the link, there are k racing exponential processes. Using the memoryless property[26] and Equation (2), the expected departure time of each aircraft (relative to the present time) is given by $\eta_l + \frac{kX_l}{\mu_l}$. Consequently, the rate of each process is the inverse of this quantity, and the net departure rate for k independent exponential processes is

$$R_{kl} = k \cdot \frac{1}{\eta_l + \frac{kX_l}{\mu_l}} \text{ aircraft per unit time.} \tag{9}$$

Since each departure from the link corresponds to an expected taxi time reduction of $\frac{X_l}{\mu_l}$ for the aircraft being assigned pushback delay, the instantaneous rate of decrease of expected taxi time is given by

$$-\frac{d}{dt}t_t = R_{kl}\frac{X_l}{\mu_l}.$$

B.2 Relation between target traffic level and control weight

In order to target a certain level of traffic (say $k = k_{\text{ctrl}}$), the value of the weight is chosen such that the rate of reduction of expected taxi time is equal to the rate of increase of the term $a \cdot t_p$, when $k = k_{\text{ctrl}}$. Since the term $a \cdot t_p$ increases at the constant rate a, the relationship between the weighting factor a and the target traffic level k_{ctrl} is

$$a = -\frac{d}{dt}t_t = R_{k_{\text{ctrl}}l}\frac{X_l}{\mu_l} = \frac{k_{\text{ctrl}}\frac{X_l}{\mu_l}}{\eta_l + k_{\text{ctrl}}\frac{X_l}{\mu_l}}. \tag{10}$$

C. Optimal control strategy

Since no new aircraft enter the link until the current aircraft pushes back, R_{kl} is the rate of decrease of the projected traffic level k_p. It is therefore possible to find the optimal control t_p required to maintain a specific value of k_p. An implicit assumption is that the exponential nature of the departure process for each aircraft is maintained as this projected value evolves. If $k_p(t_p)$ is the projected traffic level after time t_p,

$$\frac{d}{dt}k_p(t_p) = -R_{k_p l} = -\frac{k_p}{\eta_l + k_p\frac{X_l}{\mu_l}}$$
$$\Rightarrow t_p = \eta_l \ln\frac{k_p(0)}{k_p(t_p)} + (k_p(0) - k_p(t_p))\frac{X_l}{\mu_l}. \tag{11}$$

Since $k_p(0)$ is a known quantity (the current traffic level), the optimal control for each $k = k_p(0)$ is defined by substituting $k_p(t_p) = k_{\text{ctrl}}$ in Equation (11). If the optimal value is negative, t_p is assigned a value of zero, in order to obey the constraint $t_p \geq 0$. Note that this happens if and only if $k_p(0) < k_{\text{ctrl}}$, which means that the control strategy calls for immediate pushback if the traffic level is below the target value. If the current traffic level is above k_{ctrl}, the

pushback delay becomes progressively larger with $k_p(0)$. Another point to note is that for every value of $k = k_p(0)$, there is a unique control t_p that is commanded. Thus there is a discrete (but infinite, since the possible values of $k_p(0)$ are infinite) set of commanded control values that is generated by this control strategy. This fact is used in the next section, when considering the regulation characteristics of the control strategy described above.

IV.   Simulations

In this section, the control strategy developed in Section III is demonstrated through simulations, first for a single-link network, and then for the full surface network of Boston Logan.

A. Single-link example

Since the optimal control strategy as derived above relies on a number of approximations, it is necessary to validate it using independent simulations. In Figure 7 and Figure 8, simulation results are shown for a single-link network, set up using the model from Equation (1) and the controls defined by Equation (11). It is assumed that there is an infinite buffer of aircraft waiting for release, and that the current traffic level is known at all times. In addition, aircraft are released according to a First-Come-First-Served (FCFS) policy. From Figure 7, the average steady-state traffic level is seen to stabilize to k_{ctrl}. Figure 8 shows that the average taxi times are close to the value of 105.5 sec predicted by Equation (2). As predicted by Equations (6) and (7), $k_{\mathrm{ctrl}} = 10$ adds on average, $\frac{\eta_l}{k_{\mathrm{ctrl}}} = 6$ sec to the minimum inter-departure times. However, this throughput shortfall can made arbitrarily close to zero by increasing k_{ctrl}.
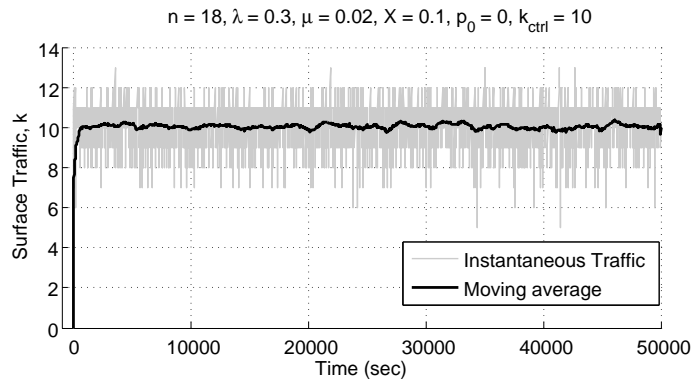

Figure 7. **Variation of surface traffic levels seen in simulations of control algorithm for single link, with** $k_{\mathrm{ctrl}} = 10$
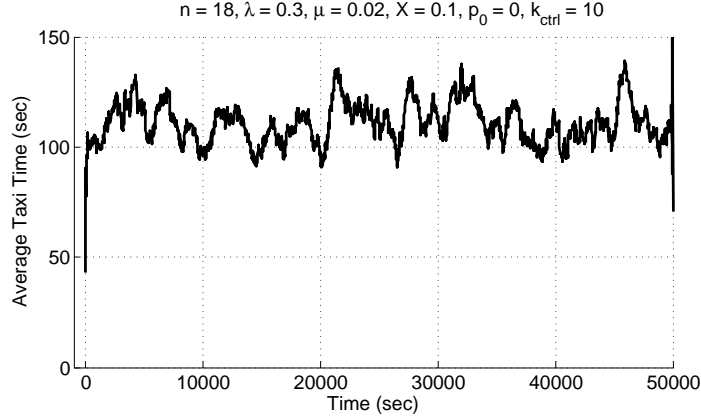
Figure 8. Variation of average taxi-out times seen in simulations of control algorithm for single link, with $k_{\mathrm{ctrl}} = 10$

B. Complete Boston Logan Airport network example

Figure 9 and Figure 10 show results from a simulation based on the full network abstraction of Boston Logan airport. All departures are assumed to happen from Runway 27 (marked as node 6 in Figure 1). The relevant network is depicted in Figure 2. It is assumed that there is a finite number of aircraft waiting for release at any time. The pushback requests from these aircraft appear at a rate based on the historical variation over a day, as derived from surface surveillance data. The x-axis shows the local time at the airport. Each curve in Figure 9 (a) plots the variation of simulated *average* traffic level in that 15-minute interval over the course of the day, for different values of k_{ctrl}. In all the restricted cases, the control strategy successfully limits surface traffic levels to the corresponding target values in times of high departure demand. Since the same pushback schedule is used for all four tests, the total number of departures is the same when summed over the whole day. The curve for unrestricted entry (or $k_{\mathrm{ctrl}} = \infty$) into the network clearly shows the morning and evening demand peaks, with high levels of surface congestion during these times. The cases with $k_{\mathrm{ctrl}} = 15$ and $k_{\mathrm{ctrl}} = 20$ mitigate this effect to a large extent, by delaying aircraft at the gate during periods of excessive demand. Note that the traffic peaks become wider with decreasing k_{ctrl}, as the control algorithm takes longer to clear out delayed aircraft. For $k_{\mathrm{ctrl}} = 10$, the peaks are so wide that they merge with each other, and the algorithm is unable to clear the built-up demand until the end of the day. This value of k_{ctrl} is too aggressive, as seen from the excessive pushback delays assigned in this case (Figure 10 (b)).

Practically, the choice of k_{ctrl} may be made on a case-to-case basis for each airport. It should be low enough to avoid gridlock on the surface, which leads to large and highly variable delays. At the same time, it should be high enough to keep up with pushback demand and the throughput capacity of the airport. For example, in the unrestricted case, a comparison of Figure 9 (a) and Figure 9 (b) shows that the peak traffic level is higher than the maximum departure rate from the airport, leading to a buildup of traffic. On the other hand, the maximum departure rate of

around 20 aircraft per 15 min is never achieved by setting $k_{\text{ctrl}} = 10$. This is because the surface traffic density is too sparse for 20 aircraft to reach the runway within 15 min. Once a feasible range of k_{ctrl} is derived from such considerations, the value can be tuned by deciding on the relative importance of one minute of taxi time reduction versus one minute of pushback delay.

The parameter a in the objective function (Equation (8)) determines the relative weighting of taxi-out time reduction and pushback delay, assuming small perturbations from the initial state. Its value varies slightly depending on an aircraft's source node, since the expected taxi times are slightly different. Figure 10 (b) shows the average taxi-out times per 15-min interval by time of day. We notice that during peak demand times, taxi-out times for the unrestricted case can be as high as 24 min. By contrast, a value of $k_{\text{ctrl}} = 10$ (which corresponds to an average value of $a = 0.45$) results in average taxi-out times of around 12 min. Similarly, when $k_{\text{ctrl}} = 15$ (average $a = 0.55$), the algorithm regulates the system to the point where 1 minute of additional pushback delay is traded off for 0.55 min, or 33 sec, of taxi time reduction.
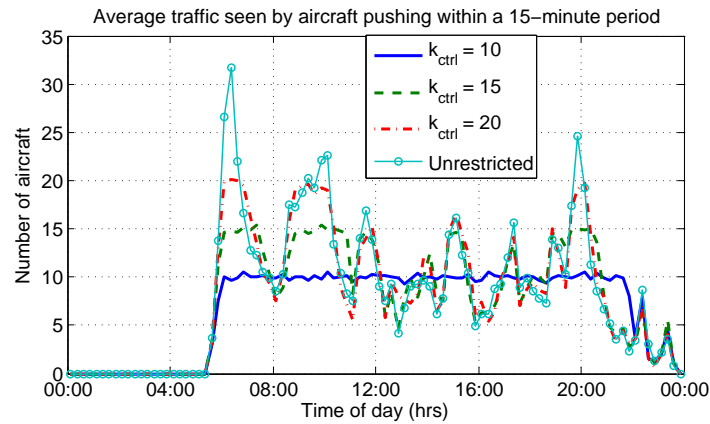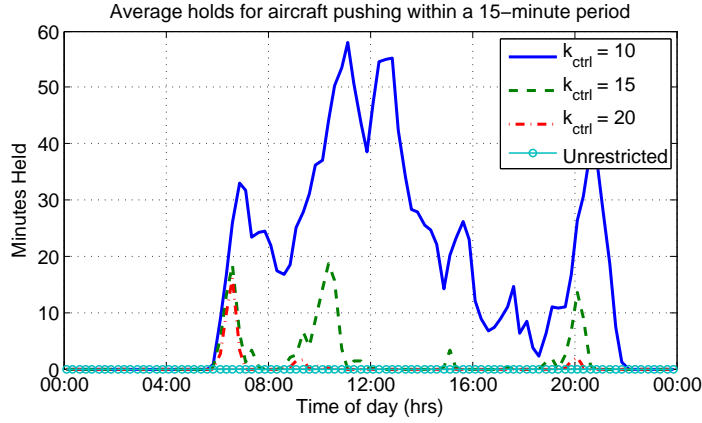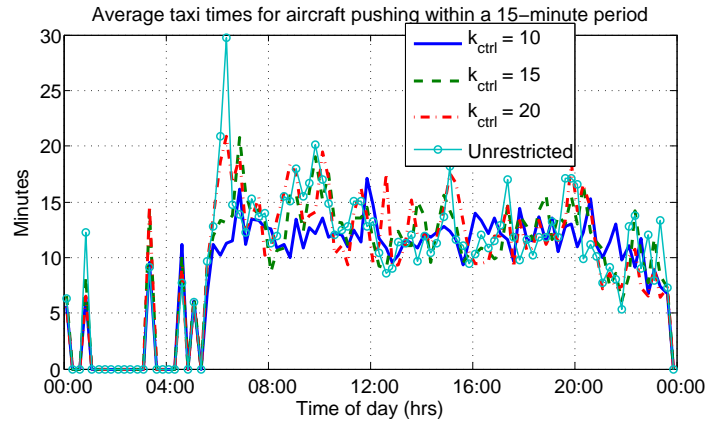

(a) Average traffic levels
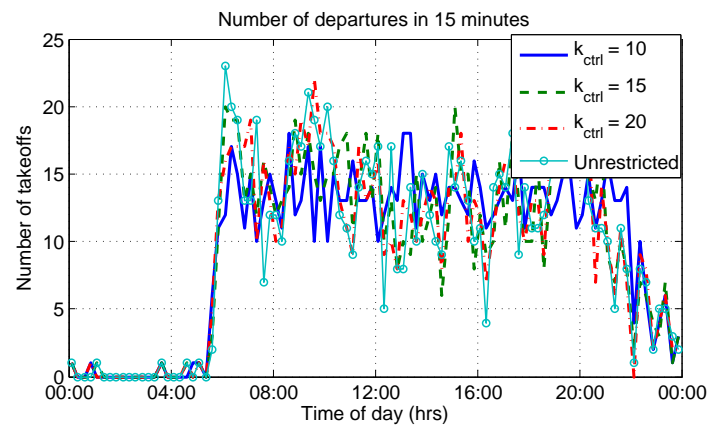
(b) Throughput per 15 minutes

Figure 9. Simulated traffic levels and throughput for the Boston Logan network, with departures from Runway 27

(a) Average gate delays

(b) Average taxi-out times

Figure 10. Simulated gate delays and taxi times for the Boston Logan network, with departures from Runway 27

C. Regulation characteristics of proposed control algorithm

C.1 Markov chain model

To analyze the effect of the control algorithm on traffic level, it is useful to consider the single-link case. As noted earlier, a fixed delay is assigned to each state, based on the target traffic level. Since departures from the link are stochastic, the state seen by the next aircraft to pushback is a random variable. It is possible to calculate the probability mass function of this state conditioned upon the previous state and the delay assigned. A Markov chain structure is seen to emerge, with the states corresponding to the traffic level seen by successive aircraft, and with known transition probabilities out of each state. Figures 11 and 12 depict this structure, assuming that the target traffic level is $k_{\text{ctrl}} = k - 1$. For states at or below k_{ctrl}, no delay is assigned according to Equation (11), and the transition is deterministic. For states above k_{ctrl}, the transitions are stochastic.
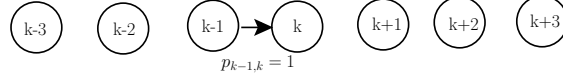
Figure 11. Markov chain model for transitions under the proposed control strategy, if $k < k_{\text{ctrl}}$
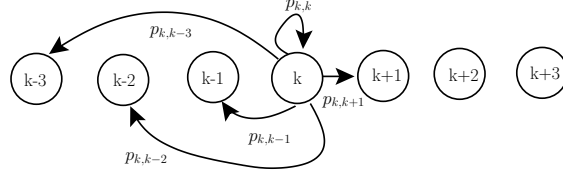

Figure 12. The Markov chain model for transitions under the proposed control strategy, if $k \geq k_{\text{ctrl}}$

C.2 Calculation of transition probabilities

Consider the scenario depicted in Figure 12, in which the current aircraft sees a traffic level of k. Then, according to Equation (11), it is assigned a pushback delay

$$t_p = \eta_l \ln \frac{k}{k-1} + (k - (k-1)) \frac{X_l}{\mu_l}.$$

From Equation (9), the Poisson departure rate from the link is given by $R_{kl} = \frac{k}{\eta_l + \frac{kX_l}{\mu_l}}$ aircraft per unit time. Note that a transition from state k to $k+1$ will occur if no aircraft depart from the link in time t_p. From the Poisson distribution, the probability of this transition is

$$p_{k,k+1} = \frac{(R_{kl}t_p)^0 \, e^{-R_{kl}t_p}}{0!} = e^{-R_{kl}t_p}.$$

In a similar manner, the transition probabilities $p_{i,j}$ can be calculated for each pair $\{i, j\}$. The chain may be truncated at a state K sufficiently larger than k_{ctrl}, such that the probability of ever reaching state K is very low. Steady state occupation probabilities for $k = 0, 1, \ldots, K$ are then given by the left eigenvector of the matrix of transition probabilities $p_{i,j}$. This matrix is nearly stochastic for a sufficiently large value of K. For the given system, these steady state values correspond to the proportion of traffic states seen by successive aircraft, when released according to the control algorithm defined by Equation (11). As shown in Figure 13 for $k_{\text{ctrl}} = 10$, these theoretical values match well with fractional state occupancies observed during simulation runs. The system is found to occupy one of the states k_{ctrl}, $k_{\text{ctrl}} \pm 1$ and $k_{\text{ctrl}} \pm 2$ more than 95% of the time. Furthermore, the same result can be easily extended to more complicated networks. In this manner, the control algorithm is able to regulate the traffic level to the desired target value.
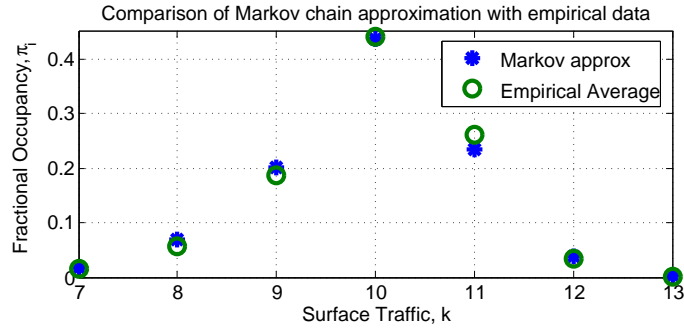
Figure 13. **Comparison of Markov-approximated steady state occupancy and average simulated occupancy, for** $k_{\mathrm{ctrl}} = 10$

V. Discussion of approach

The simulations in Section IV showed that the proposed control algorithm is capable of maintaining a steady traffic level on the airport surface under realistic conditions. Since the approach is not limited by the large time constants of algorithms previously available in literature, it reduces the variability in traffic levels seen in these studies. However, there is potential to further refine the control strategy. Specifically, the primary objective of congestion control methods is to reduce the fuel consumption and emissions on the surface. While maintaining steady, reasonable traffic levels does this to some extent, even greater savings may be possible by optimizing for these objectives directly. Therefore, there is incentive to develop control strategies that explicitly aim to minimize aircraft fuel consumption and emissions by considering the behavior of aircraft and their engines.[27]

Finally, the applicability of any proposed strategy in the current air traffic control framework is an essential consideration. Implementation of the algorithm proposed in this paper would require a continuous knowledge of current traffic levels on the airport surface. Since real-time ASDE-X feeds are already available at most major airports within the US, this is a feasible requirement. Another issue that needs to be addressed is that of fairness, in terms of the order in which aircraft depart as compared to the order in which they call ready for pushback. Although the FCFS policy is perceived to be the most equitable by airlines and air traffic controllers, it may not result in the best possible performance, in terms of fuel savings and departure throughput. Therefore, a balance needs to be struck between the performance objectives and the perceptions of the stakeholders in the system. Control algorithms with a specific focus on handling these requirements are currently under development.[28]

VI.   Conclusions

This paper modeled the aircraft taxi-out process by representing the airport surface as a network and the total taxi time as the sum of the travel times along different links in the surface trajectory of an aircraft. A set of suitable random processes was proposed to model the distribution of link travel times. A control algorithm that attempts to minimize the delay cost for aircraft while limiting the effect on throughput was developed. Using the test cases of a simple, single-link model as well as the complete airport network for BOS, it was shown that the algorithm can successfully maintain the level of surface traffic at a specified average value. When suitably calibrated, this steady traffic level can provide significant reduction in aircraft surface fuel burn without adversely affecting airport performance.

References

[1] Federal Aviation Organization, "Aviation System Performance Metrics database," `http://aspm.faa.gov/main/aspm.asp`, Retrieved December 2010.

[2] Feron, E., Hansman, R. J., Odoni, A. R., Cots, R., Delcaire, B., Hall, W., Idris, H., Muharremoglu, A., and Pujet, N., "The Departure Planner: A conceptual discussion," White paper, MIT, International Center for Air Transportation, December 1997.

[3] Simaiakis, I. and Balakrishnan, H., "Queuing models of airport departure processes for emissions reduction," *AIAA Guidance, Navigation and Control Conference and Exhibit*, Chicago, IL, August 2009.

[4] Idris, H., Clarke, J.-P., Bhuva, R., and Kang, L., "Queueing model for taxi-out time estimation," *Air Traffic Control Quarterly*, Vol. 10, No. 1, January 2002, pp. 1–22.

[5] Carr, F., Evans, A., Clarke, J.-P., and Feron, E., "Modeling and control of airport queueing dynamics under severe flow restrictions," *American Control Conference*, Anchorage, AL, May 2002.

[6] Brinton, C., Krozel, J., Capozzi, B., and Atkins, S., "Improved taxi prediction algorithms for the Surface Management System," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Monterey, CA, August 2002.

[7] Balakrishnan, H. and Jung, Y., "A framework for coordinated surface operations planning at Dallas-Fort Worth International Airport," *AIAA Guidance, Navigation, and Control Conference*, Hilton Head, NC, August 2007.

[8] Menon, P. K., Sweriduk, G. D., and Bilimoria, K. D., "New approach for modeling, analysis, and control of air traffic flow," *AIAA Journal of Guidance, Control and Dynamics*, Vol. 27, No. 5, September-October 2004, pp. 737–744.

[9] Menon, P. K., Sweriduk, G. D., Lam, T., Diaz, G. M., and Bilimoria, K. D., "Computer-aided Eulerian air traffic flow modeling and predictive control," *AIAA Journal of Guidance, Control and Dynamics*, Vol. 29, No. 1, January-February 2006, pp. 12–19.

[10] Sensis Corporation, East Syracuse, NY, *ASDE-X brochure - ASDE-X 10/05.qxd*, 2008.

[11] Eryilmaz, A. and Srikant, R., "Fair resource allocation in wireless networks using queue-length-based scheduling and congestion control," *INFOCOM 2005, 24th Annual Joint Conference of the IEEE Computer and Communications Societies*, Miami, FL, March 2005.

[12] Neely, M., Modiano, E., and Li, C.-P., "Fairness and optimal stochastic control for heterogeneous networks," *IEEE/ACM Transactions on Networking*, Vol. 16, No. 2, April 2008, pp. 396–409.

[13] Paganini, F., Doyle, J., and Low, S., "Scalable laws for stable network congestion control," *40th IEEE Conference on Decision and Control*, Orlando, FL, December 2001.

[14] Low, S., Paganini, F., Wang, J., Adlakha, S., and Doyle, J., "Dynamics of TCP/RED and a scalable control," *INFOCOM 2002, 21st Annual Joint Conference of the IEEE Computer and Communications Societies*, New York, NY, June 2002.

[15] Perkins, J. and Kumar, P., "Stable, distributed, real-time scheduling of flexible manufacturing/assembly/disassembly systems," *IEEE Transactions on Automatic Control*, Vol. 34, No. 2, February 1989, pp. 139–148.

[16] Kumar, P. and Seidman, T., "Dynamic instabilities and stabilization methods in distributed real-time scheduling of manufacturing systems," *IEEE Transactions on Automatic Control*, Vol. 35, No. 3, March 1990, pp. 289–298.

[17] Gershwin, S. B., "Design and operation of manufacturing systems: the control-point policy," *IIE Transactions*, Vol. 32, 2000, pp. 891–906.

[18] Nikolova, E., Brand, M., and Karger, D., "Optimal route planning under uncertainty," *International Conference on Automated Planning and Scheduling*, Lake District, England, June 2006.

[19] Nikolova, E. and Karger, D., "Route planning under uncertainty: The Canadian traveller problem," *23rd Conference on Artificial Intelligence*, Chicago, IL, July 2008.

[20] Simaiakis, I., Khadilkar, H., Balakrishnan, H., Reynolds, T. G., Hansman, R. J., Reilly, B., and Urlass, S., "Demonstration of reduced airport congestion through pushback rate control," *9th USA/Europe Air Traffic Management Research and Development Seminar*, Berlin, Germany, June 2011.

[21] Zheng, Q., Zhao, Y. J., and Capozzi, B., "Time-of-arrival taxi conformance monitoring for surface operations," *AIAA Journal of Guidance, Control and Dynamics*, Vol. 34, No. 3, May-June 2011, pp. 750–760.

[22] Khadilkar, H. and Balakrishnan, H., "A multi-modal unscented Kalman filter for inference of aircraft position and taxi mode from surface surveillance data," *AIAA Conference on Aviation Technology, Integration and Operations*, Virginia Beach, VA, September 2011.

[23] Kullback, S. and Leibler, R., "On information and sufficiency," *Annals of Mathematical Statistics*, Vol. 22, No. 1, 1951, pp. 79–86.

[24] Ford, L. and Fulkerson, D., "Maximal flow through a network," *Canadian Journal of Mathematics*, 1956, pp. 399–404.

[25] Elias, P., Feinstein, A., and Shannon, C., "A note on the maximum flow through a network," *IRE Transactions on Information Theory*, Vol. 2, No. 4, December 1956, pp. 117–119.

[26] Gallager, R., *Discrete stochastic processes*, Kluwer Academic Publishers, 1996, Chapter 2: Poisson Processes.

[27] Khadilkar, H. and Balakrishnan, H., "Estimation of aircraft taxi-out fuel burn using Flight Data Recorder archives," *AIAA Guidance, Navigation, and Control Conference*, Portland, OR, August 2011.

[28] Khadilkar, H. and Balakrishnan, H., "Control of aircraft pushbacks at an airport using a dynamic programming formulation," *AIAA Guidance, Navigation, and Control Conference*, Minneapolis, MN, August 2012.