

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Impact of Off-Block Time Uncertainty on the Control of Airport Surface Operations

Sandeep Badrinath and Hamsa Balakrishnan

Massachusetts Institute of Technology, Cambridge, MA 02139, USA.

Emily Joback and Tom G. Reynolds

MIT Lincoln Laboratory, Lexington, MA 02421, USA.

Congestion at major airports worldwide results in increased taxi times, fuel burn and emissions. Regulating the pushback of aircraft from their gates, also known as departure metering, is a promising approach to mitigating surface congestion. Departure metering algorithms require models of airport surface traffic, and knowledge of when a flight would be to be ready for pushback, called the Earliest Off-Block Time (EOBT). While EOBTs are known to be inaccurate due to several reasons, there has been little prior research on characterizing EOBT uncertainty and its impact on departure metering.

We present a new class of queuing network models for the airport surface that are capable of capturing congestion at multiple locations. We demonstrate our modeling approach using operational data from three major US airports: Newark Liberty International airport (EWR), Dallas Fort Worth International airport (DFW), and Charlotte Douglas International airport (CLT). We analyze the current levels of uncertainty in the EOBT information published by the airlines and conduct a parametric analysis of the reduction in departure metering benefits due to errors in the EOBT information. Our analysis indicates that the current levels of EOBT uncertainty lead to a 50% reduction in benefits at some airports when compared to an ideal case with no EOBT uncertainty. Two approaches to departure metering are considered: NASA's ATD-2 logic, and a new optimal control approach. We show that our queuing network models can help design and evaluate both approaches, and that the optimal control approach is more effective in accommodating EOBT uncertainty while maintaining runway utilization.

Key words: Airport surface operations; Taxi-out processes; Departure metering; Earliest Off-Block Time; Uncertainty

History:

1. Introduction

The imbalance between capacity and demand, particularly during periods of peak traffic, has led to surface congestion at major airports worldwide. Delays resulting from airport congestion have an

economic impact on airlines and passengers, as well as an environmental impact due to excessive fuel consumption. Congestion-related delays at an airport can also propagate due to the networked nature of the air transportation system (Pyrgiotis, Malone, and Odoni 2013, Gopalakrishnan, Balakrishnan, and Jordan 2016).

Surface congestion mitigation approaches fall broadly into two categories: Demand management (either strategic approaches such as slot constraints, or tactical approaches such as departure metering) and capacity improvements (Cheng et al. 2016, Diana 2018, Morrison and Winston 2007, Ryerson and Woodburn 2014). The latter class of approaches tends to be prohibitively expensive, requiring significant changes in policy or additional infrastructure. By contrast, tactical demand management approaches such as departure metering have been shown to reduce taxi-out delays (Nakahara et al. 2011, Simaiakis et al. 2014, Sharma et al. 2018). Departure metering at the top 35 airports in the US has been estimated to have potential fuel savings of 2.2-3.9 billion gallons over a 20-year period, reducing costs by \$5.5-9.5 billion (Nakahara and Reynolds 2013).

Several approaches to departure metering, based both on heuristics and optimization algorithms, have been previously investigated (Pujet, Delcaire, and Feron 2000, Verma et al. 2018, Simaiakis, Sandberg, and Balakrishnan 2014, Khadilkar and Balakrishnan 2014). The underlying objective of departure metering is to assign appropriate holds for departing aircraft at their gates, so as to minimize taxi-out times while maintaining runway utilization. The taxi-out process is subject to significant uncertainties, including the time that an aircraft is ready to pushback, its pushback time, taxi routes, taxi speeds, and separation between successive take-offs (Rappaport et al. 2009, Lee 2014). Prior work on departure metering under uncertainty has generally focused on the uncertainty in taxi times and runway capacity (Burgain et al. 2012, Simaiakis, Sandberg, and Balakrishnan 2014, Bosson and Sun 2016, Murça 2017). By contrast, demand-side uncertainty, namely, the uncertainty associated with the time at which an aircraft will be ready to pushback, has received little attention.

The predicted departure demand a few hours prior to operations is based on the Scheduled Off-Block Times (SOBTs) published by the airlines. Closer to the time of operation, departure demand can be based on the Earliest Off-Block Times (EOBTs), which represent the times at which flights are expected to be ready for pushback. *Note: We adopt the nomenclature used by the US Surface Collaborative Decision Making (S-CDM) community for different data elements. The nomenclature used by EUROCONTROL's Airport-CDM is significantly different; for example, EOBT is referred to as Estimated Off-Block Time (FAA Surface CDM Team 2012, EUROCONTROL Airport CDM Team 2017).* In the US, airlines are now encouraged to publish EOBT information for their flights, and they are expected to dynamically update each of them until the flight pushes back from the gate (FAA Surface CDM Team 2012). The dynamic updates are intended to make the EOBT more

representative of the actual demand compared to the static SOBT. Most departure metering concepts assume the availability of EOBT information while making tactical gate-hold decisions (FAA Surface CDM Team 2012, EUROCONTROL Airport CDM Team 2017). While prior efforts have focused on near real-time departure metering (for example, over the next 15 minutes), departure metering concepts call for assigning the gate-hold times over longer time-horizons, in order to improve the predictability of operations (Liu et al. 2014). The EOBT is also a key input to the calculation of downstream events within the S-CDM concept (Okuniek and Sparenberg 2017). The actual pushback time (also referred to as the *call-ready* or *push-ready* time) of a flight often differs from the published EOBT due to several reasons. The accuracy of the published EOBT depends on an airline’s operating policies with respect to connecting passengers, and aircraft and crew schedules, as well as CDM procedures (FAA Surface CDM Team 2012, Ball et al. 2001). Despite the significant influence of EOBT accuracy on surface operations, there has, to the best of our knowledge, been limited research on characterizing its magnitude and impact of the resulting uncertainty on the efficacy of departure metering (Ball, Vossen, and Hoffman 2001, McFarlane and Balakrishnan 2016).

Departure metering algorithms generally require predictions of taxi-out times, which can be obtained using a number of approaches including queuing models, stochastic node-link models, discrete-event simulations, and models based on statistical regression (Simaiakis and Pyrgiotis 2010, Khadilkar 2013, Windhorst et al. 2013, Lee et al. 2015). Queuing models have been widely used to model the taxi-out process since queues naturally reflect the congestion that occurs at an airport (Simaiakis and Balakrishnan 2015, Jacquillat 2012). Most often, a single departure runway queue is used to represent congestion, and the taxi-out time is given by the sum of the unimpeded time to the runway and wait time in the queue (Simaiakis and Balakrishnan 2015). However, more complex queuing network models may be needed for airports that have several points of congestion or multiple departure runways (Badrinath, Li, and Balakrishnan 2018).

In this paper, we present an approach to developing queuing models of airport surface operations that can be easily adapted to a range of layouts and operating conditions. After validating these models for three major US airports, we use them to conduct a parametric study of the impact of EOBT uncertainty on two approaches to departure metering: NASA’s Airspace Technology Demonstration-2 (ATD-2) logic (Verma et al. 2018) and a new optimal control approach. We also analyze the EOBT information published by airlines at these airports in order to evaluate the current levels of uncertainty. The three airports considered – Dallas Fort-Worth International Airport (DFW), Charlotte Douglas International Airport (CLT) and Newark Liberty International Airport (EWR) – are chosen in part because of the differences in layout and traffic. DFW (621,684 aircraft movements in 2017, the 4th busiest airport in the world by number of movements) and CLT (546,845 aircraft movements, the 7th busiest airport in the world by number of movements)

are major hubs for American Airlines (Federal Aviation Administration 2018b). There have been demonstrations of departure metering at CLT under Phase 2 of the ATD-2 program since November 2017, and DFW has been selected as the site for Phase 3 of the program. EWR handled 432,941 operations in 2017, is a hub for United airlines in the highly-constrained New York metroplex, and is prone to frequent congestion.

The models and analyses presented in this paper are of potential benefit to multiple stakeholders such as airlines, airport operators, and air navigation service providers (such as the FAA or EUROCONTROL). For example, emerging terminal automation systems such as the FAA Terminal Flight Data Manager (TFDM) rely on accurate EOBTs to be published for input to their departure metering functions. Improvements in EOBT accuracy will require investments on the part of the airlines in terms of new technologies and procedures. Our analyses help estimate the returns of these investments in terms of taxi-out time and fuel burn savings. We show that the proposed class of queuing network models can be easily adapted to new airports and operating environments, and can help improve predictability as well as environmental impact assessments (Liu et al. 2014, Nakahara and Reynolds 2013).

2. Queuing models of airport surface operations

Queuing models can be used to determine the lengths of various queues on the airport surface in order to predict the taxi-out times of departures. The taxi-out time of a flight is determined as the sum of the unimpeded travel time (the time it would take for the aircraft to move from the gate to the runway with no congestion) and the wait time in queues. In general, queues are formed when the demand for resources exceeds the available capacity. In the case of the airport surface, runways and certain taxiways become bottlenecks during periods of heavy traffic, leading to the formation of queues. Since the level of congestion depends on demand, aircraft fleet mix, airport procedures, taxi routes, runway configuration (set of runways used for arrivals and/or departures) and weather, queuing models need to be adapted to different operating conditions.

2.1. Queuing network models of EWR, DFW and CLT

Figure 1 shows the airport layouts of EWR, DFW and CLT, along with a snapshot of the traffic movements on the surface to illustrate surface queues. Departures and arrivals are represented as black and white triangles, respectively. Each of these airports operates primarily in one of two runway configurations: North Flow (NF) and South Flow (SF). Figure 1 illustrates the most prevalent runway configuration: SF for EWR (58% of the time in 2017), SF for DFW (70% of the time) and NF for CLT (53% of the time). The layout for DFW is cropped to show only the primary runways used in South Flow.

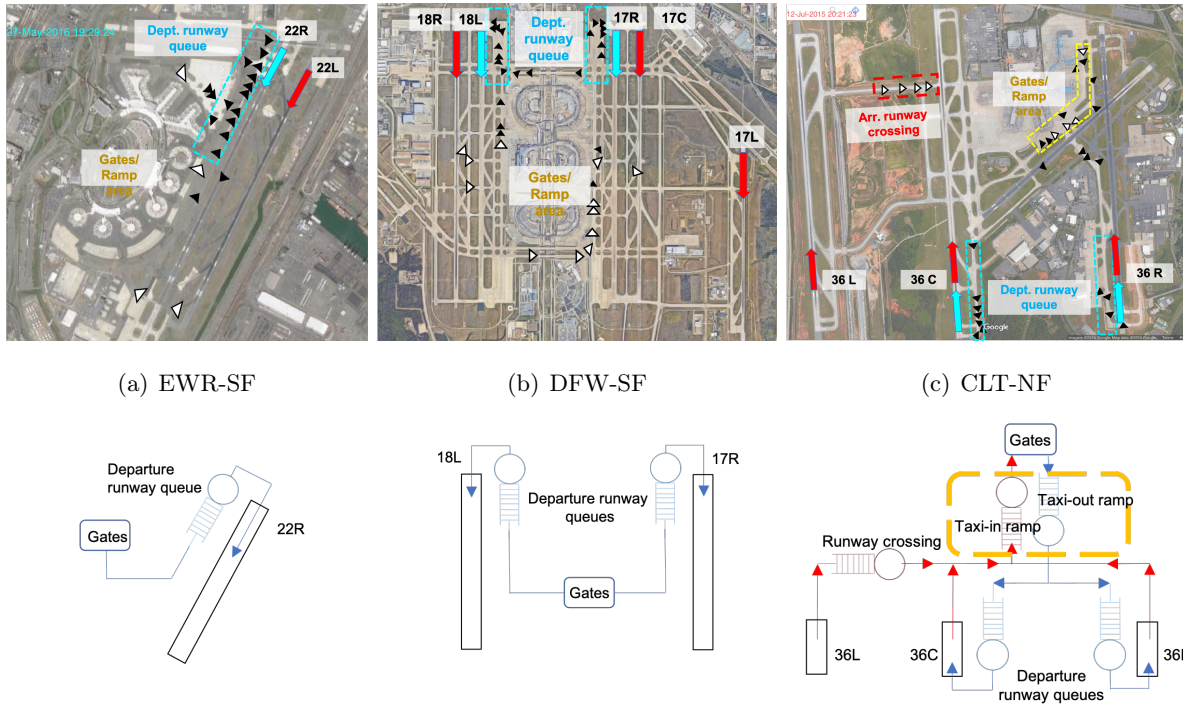


Figure 1 Snapshots of traffic movements at EWR, DFW and CLT, and the corresponding queuing representations under the most frequently-used configurations. The primary departure and arrival runways are indicated in the figure using red and blue arrows, respectively. Departures and arrivals are represented as black and white triangles, respectively.

All three airports experience large queues of aircraft waiting for take-off near the departure runways. The number of departure runway queues varies by airport and runway configuration: the scenarios considered in this paper result in one departure runway queue at EWR (22R), and two each at DFW (18L, 17R) and CLT (36C, 36R). Runway separation requirements are the primary driver of their capacity, and depend on weather (instrument vs. visual meteorological conditions, IMC or VMC), the fleet mix of aircraft using the runway, and the relative proportion of arrivals and departures (Simaiakis and Balakrishnan 2014).

Some airports (for example, CLT) experience queuing in the ramp area (region close to the airport terminal building) in addition to near the departure runways. Consequently, flights at CLT spend nearly half of their taxi-out time in the ramp area (Table 1). Queuing network models can be used to represent multiple points of congestion on the airport surface. Figure 1 shows examples of queuing network models for EWR, DFW and CLT. Since EWR and DFW do not exhibit significant ramp congestion, queuing is assumed to occur primarily at the departure runway(s). By contrast, the queuing network model for CLT includes a ramp queue in addition to two departure runway queues. Arrivals and departures interact in the ramp area, motivating the need to model the taxi-in process. The taxi-in process in CLT-NF is represented as follows: flights landing on the leftmost runway

(36L) pass through a runway crossing queue and a taxi-in ramp queue, whereas flights landing on one of the other runways just pass through the taxi-in ramp queue. The service rate of the taxi-out ramp server is modeled as a function of the taxi-in ramp queue length, and vice versa (Badrinath, Li, and Balakrishnan 2018). The three queuing network models (EWR-SF, DFW-SF and CLT-NF) correspond to a progression of complexity, going from a single runway queue to two runway queues, and finally multiple runway and ramp queues.

Airport	Config.	Avg. gate-to-spot (min)			Avg. spot-to-runway (min)			Avg. taxi-out (min)
		Unimpeded	Delay	Total	Unimpeded	Delay	Total	
EWR	NF	5.7	3.1	8.8	6.0	6.4	12.4	21.2
	SF	5.9	3.2	9.1	5.4	6.1	11.4	20.6
DFW	NF	5.1	2.4	7.4	6.5	4.4	10.9	18.4
	SF	5.1	2.4	7.5	5.4	3.9	9.4	16.8
CLT	NF	6.3	3.5	9.8	5.1	5.3	10.4	20.2
	SF	7.4	4.2	11.6	3.2	4.4	7.6	19.2

Table 1 Taxi-out times, including the unimpeded and delay components, at EWR, DFW and CLT, separated by ramp area (gate-to-spot) and active movement area (spot-to-runway). The statistics are computed using multiple data sources (OAG 2016, Federal Aviation Administration 2018a,b) for a three month period (May-July 2016) for DFW and EWR, and a five month period for CLT (May-July 2015; May-Jun 2016).

The taxi-out processes under each of the secondary runway configurations (EWR-NF, DFW-NF, CLT-SF) is represented using a single queue for each active departure runway. It is worth noting that while the ramp area taxi-out delays in Table 1 are not negligible for EWR (NF and SF) and CLT (SF), this is in part due to an overflow of the runway queue into the ramp area.

2.2. Models of queue dynamics




Figure 2 Number of pushbacks (per 15 min), declared departure capacity (per 15 min) and total runway queue length for a typical day at CLT (05/22/2016).

A distinguishing feature of queues on the airport surface is their non-stationarity, which is driven by the time-varying nature of demand and capacity. Figure 2 shows the number of pushbacks,

declared departure capacity (Federal Aviation Administration 2018b), and total runway queue length at CLT on a representative good weather day (5/22/2016). The variations in departure demand (pushbacks) are as expected at a hub airport with banked operations. Demand exceeds departure capacity during each departure bank, leading to the formation of large queues that fluctuate over the course of the day. Capacity also varies because of changes in weather and arrival rate, resulting in time-varying queues.

2.2.1. Fluid model of non-stationary queues. The exact analysis of non-stationary queuing networks is analytically challenging, motivating the use of numerical simulations or approximations (Kivestu 1976, Hansen et al. 2009, Lovell et al. 2013). Probabilistic models using Markov chains can be restrictive in terms of the service time distributions they can accommodate, and lead to high-dimensional state spaces for large queuing networks. The size of the state space also poses a challenge to the development of control strategies. A key innovation of this work is the development of a fluid model of the time-varying taxi-out queuing network. The model is a continuous approximation to the discrete queuing problem, derived by combining results from steady-state queuing theory with the flow conservation principle, to obtain a point-wise stationary approximation (Tipper and Sundareshan 1990, Wang, Tipper, and Banerjee 1996, Badrinath and Balakrishnan 2017).

2.2.2. Single queue and server. We first consider the simple case of a single queue and a single server, with time-varying arrival and service rates. Let $x(t)$ represent the average number of customers in the queue at time t . Let $\lambda(t)$ and $\mu(t)$ denote the average arrival rate and service rate at time t , respectively. The word "average" here denotes the ensemble average at a particular time-instant. Let $f_{\text{in}}(t)$ and $f_{\text{out}}(t)$ represent the in-flow and out-flow rate from the queue at time t . From the flow conservation principle, we have:

$$\dot{x}(t) = -f_{\text{out}}(t) + f_{\text{in}}(t). \quad (1)$$

Assuming that there are no constraints on the queue length, $f_{\text{in}}(t) = \lambda(t)$. For the out-flow, $f_{\text{out}}(t) = \mu(t)\rho(t)$, where $\rho(t)$ is the average utilization of the server. The queue dynamics is given by:

$$\dot{x}(t) = -\mu(t)\rho(t) + \lambda(t). \quad (2)$$

The average utilization, $\rho(t)$, is approximated by a function, $G(x(t))$, which satisfies the following properties: (a) $G(0) = 0$ and $G(\infty) = 1$; and (b) $G(x)$ is strictly concave and nonnegative $\forall x \in [0, \infty)$, in order to represent congestion. The dynamics for $x(t)$ can then be rewritten in terms of $G(x)$ as:

$$\dot{x}(t) = -\mu(t)G(x) + \lambda(t), \quad x(0) = x_0. \quad (3)$$

The expression for $G(x)$ is obtained by matching the steady-state number of customers in the system. Assuming a Poisson arrival process, the Pollaczek-Khinchine formula (4) provides an expression for the mean number of customers (x_s) at steady state (Wang, Tipper, and Banerjee 1996), namely,

$$x_s = \rho + \frac{\rho^2(1 + C_v^2)}{2(1 - \rho)}. \quad (4)$$

Here, C_v is the coefficient of variation of the service time distribution. Expressing ρ in terms of x_s , we get:

$$\rho = \frac{x_s + 1 - \sqrt{x_s^2 + 2C_v^2 x_s + 1}}{1 - C_v^2}. \quad (5)$$

Rewriting x_s in terms of ρ and using the fact that $G(x)$ is an approximation of ρ :

$$\rho(t) \approx G(x) = \frac{x + 1 - \sqrt{x^2 + 2C_v^2 x + 1}}{1 - C_v^2}. \quad (6)$$

For a server with exponential service time distribution, $C_v = 1$, and the above expression for $G(x)$ simplifies to $x/(1+x)$. Motivated by this expression, the function $G(x)$ is approximated by $Cx/(1+Cx)$, where the parameter C is given by:

$$\arg \min_C \int_0^{x_m} \left(G(x) - \frac{Cx}{1+Cx} \right)^2 dx, \quad (7)$$

where x_m denotes the maximum queue size expected in the system. The parameter C is essentially an empirical value that depends on C_v and x_m . The dependence of the C parameter on C_v for $x_m = 10$ is shown in Figure 3(a). Additionally, the value of C does not change significantly with x_m if it is considered to be sufficiently large (see Figure 3(b)). This is primarily because $G(x)$ and its approximation saturate close to one for large values of x . Moreover, C is less sensitive to x_m for larger values of C_v and it is completely independent of x_m for $C_v = 1$ (case with exponential service time distribution).

Figure 4 compares the approximation for $G(x)$ and the actual value for the case of an Erlang distribution with shape parameter of 10 and rate of 5, resulting in a coefficient of variation (C_v) of 0.2. In this example, if $x_m = 10$, then $C = 1.52$. There is a good match between the approximation and the actual value. Finally, the above approximation for $G(x)$ results in the following ODE for the evolution of the mean queue length:

$$\dot{x}(t) = -\mu(t) \frac{C(t)x(t)}{1 + C(t)x(t)} + \lambda(t). \quad (8)$$

We illustrate the performance of the analytical queuing model for some standard queues. Figure 5 compares the mean queue length for a $M_t/M/1$ queuing system obtained from the analytical queuing model, against the corresponding output of a discrete stochastic simulation (obtained from 3,000 independent samples). The results from the analytical queuing model are found to closely match the




Figure 3 The dependence of the C parameter on C_v and x_m .




Figure 4 Comparison between the approximation for $G(x)$ and the actual value ($C_v = 0.2, x_m = 10, C = 1.5$).




Figure 5 Comparison of the mean queue lengths obtained from simulations and the proposed analytical queuing model for an $M_t/M/1$ queue. Note the different x-axis limits.

simulation. As the load on the server ($\frac{\lambda(t)}{\mu}$) increases, the model begins to deviate from the discrete simulation (Figure 5(b)). The model also shows deviations if the queue length fluctuates rapidly, as seen in Figure 5(c), which was obtained by increasing the sinusoidal frequency of the arrival

rate into the queue. Such a deviation is expected since the assumption of a point-wise stationary approximation breaks down. Finally, although prior work has considered finite queue buffers for point-wise stationary fluid flow models (Wang, Tipper, and Banerjee 1996), the results are too complex for practical implementation. However, the limitation of unconstrained queue sizes does not preclude these analytical queuing models from being used at most major airports.

We consider the case of an Erlang service time distribution (Figure 6), which is encountered frequently in communication networks and transportation systems. Once again, the analytical model is a close match to the exact simulations, while offering a considerable computational advantage.




Figure 6 Comparison of the mean queue lengths obtained from simulations and the proposed analytical queuing model for an $M_t/E_{10}/1$ queue. Note the different x-axis limits.

2.2.3. Multiclass queues. The model for a single queue can be extended to handle multiple classes of customers (Tipper and Sundareshan 1990). Let $i = 1, 2, \dots, l$ denote different classes of customers in the system, and x_i be the number of customers of class i in the queue buffer. The evolution of total number of customers in the queue ($x_T = \sum_{i=1}^l x_i$) can be obtained using Eq. (8):

$$\dot{x}_T = -\mu \frac{Cx_T}{1 + Cx_T} + \lambda. \quad (9)$$

The effective mean service rate for each class in the queuing dynamics is assumed to be proportional to the fraction of customers of that particular class in the queue buffer, considering service time distribution is same for all the customers. Using this assumption, the evolution of mean queue length of a particular class i is given by:

$$\dot{x}_i = -\mu \frac{Cx_T}{1 + Cx_T} \frac{x_i}{x_T} + \lambda = -\mu \frac{Cx_i}{1 + Cx_T} + \lambda. \quad (10)$$

2.2.4. Queuing networks. The single queue model can be extended to a network of multiple queues using the flow conservation principle: The output of one queue becomes the input to the second queue if they are connected. Let R be the routing matrix, with elements r_{ij} representing the fraction of customers joining queue j after being served by server i . Let λ_i be the exogenous input into queue i with mean service rate μ_i . The dynamics of the mean queue length for each queue in the network is given by:

$$\dot{x}_i = -\mu_i \frac{C_i x_i}{1 + C_i x_i} + \lambda_i(t) + \sum_j \mu_j \frac{C_j x_j}{1 + C_j x_j} r_{ji}. \quad (11)$$

While deriving the analytical model for a single queue, we assumed that the arrival process into the queue is Poisson. However, the output of one queue that becomes the input to the second queue need not follow a Poisson process for a general service time distribution. We make an approximation that the arrivals into the queue follows a Poisson process even in the case of a queuing network to obtain Eq. (11). However, one needs to note that this particular approximation might not work well for queues with small buffer capacity (which is not a constraint for the current problem of interest).

2.2.5. Queuing networks with time-delays. Another aspect that arises in many realistic queuing networks are time-delays due to propagation. Note that the propagation delay does not include the wait time in the queue. Let τ_{ij} be the propagation time (travel time) from server i to j . Then, the mean queue length is given by the following delay differential equation:

$$\dot{x}_i = -\mu_i \frac{C_i x_i}{1 + C_i x_i} + \lambda_i + \sum_j \mu_j (t - \tau_{ji}) \frac{C_j (t - \tau_{ji}) x_j (t - \tau_{ji})}{1 + C_j (t - \tau_{ji}) x_j (t - \tau_{ji})} r_{ji} (t - \tau_{ji}). \quad (12)$$

Since the arrival times and service times are stochastic in nature, the queue length is a random variable. The proposed analytical queuing model governs the evolution of the ensemble mean queue length. By contrast, probabilistic queuing models such as Markov chains provide the probabilities of queue length at any time instant. However, such probabilistic models are often complex, making it difficult to model large queuing networks and to develop feedback controllers. As an alternative, one can develop congestion control strategies by assuming that the actual queue length at any time instant is a small deviation from the ensemble mean queue length.

2.3. Application of analytical queuing models to the airport surface

Suppose the airport surface is represented using a queuing network with p queues. Let $\mathbf{x}(t) \in \mathbb{R}^p$ be a vector of queue lengths on the airport surface at any time instant t . Let $\mathbf{u}_d(t) \in \mathbb{R}^q$ be a vector of the pushback rates at time t , where each component, $\mathbf{u}_{d_i}(t)$ represents the pushback rate to the i^{th} runway. The pushback rate refers to the number of departures pushing back from the gate per unit time (say, 5-min). Similarly, let $\mathbf{u}_a(t) \in \mathbb{R}^q$ be a vector of landing rates of arrivals on the runway.

Let $\mu(t) \in \mathbb{R}^p$ be a vector of the time-varying mean service rates for each of the servers, and let $\mathbf{C}(t) \in \mathbb{R}^p$ be the vector of parameters associated with the approximation of the utilization factor for each of the servers (from Eq. (7)). The input to each queue in the network is assumed to be Poisson. The evolution of the queues can then be expressed as:

$$\dot{\mathbf{x}} = \mathbf{f}\left(\mathbf{x}(t), \mathbf{x}(t - \tau_1), \dots, \mathbf{x}(t - \tau_m), \mu(t), \mu(t - \tau_1), \dots, \mu(t - \tau_m), \mathbf{C}(t), \mathbf{C}(t - \tau_1), \dots, \mathbf{C}(t - \tau_m), \mathbf{u}_d(t - \tau_{m+1}), \dots, \mathbf{u}_d(t - \tau_w), \mathbf{u}_a(t - \tau_{m+1}), \dots, \mathbf{u}_a(t - \tau_w)\right), \quad (13)$$

where τ_k is the time to move unimpeded between two points in the queuing network, for $k = 1, 2, \dots, w$. The function $f(\cdot)$ is obtained from the fluid model and depends on the connectivity of the queuing network. We consider the taxi-out process for the following cases:

- **[Case 1]** An airport represented using q departure runway queues:

$$\dot{x}_{r_i} = -\mu_{r_i}(t) \frac{C_{r_i}(t)x_{r_i}(t)}{C_{r_i}(t)x_{r_i}(t) + 1} + u_{d_i}(t - \tau_i), \quad i = 1, 2, \dots, q \quad (14)$$

where x_{r_i} represents the queue length of the i^{th} departure runway, and τ_i is the average unimpeded travel time from a gate to the i^{th} departure runway. The unimpeded time is computed as the 10th percentile of the taxi-time distribution. This model with only runway departure queues is used for EWR ($q = 1$), DFW ($q = 2$), and CLT-SF ($q = 2$).

- **[Case 2]** An airport modeled by a single ramp queue and q departure runway queues:

We model the ramp queue as a multi-class queue, the class of customers representing the runway assignment of the aircraft in the queue. The service rate for a particular class is proportional to the number of customers of that class in the queue. The queuing dynamics are then given by

$$x_s = \sum_{i=1}^q x_{s_i} \quad (15)$$

$$\dot{x}_{s_i} = -\mu_s(t) \frac{C_s(t)x_{s_i}(t)}{C_s(t)x_s(t) + 1} + u_{d_i}(t - \tau_{gs}) \quad (16)$$

$$\dot{x}_{r_i} = -\mu_{r_i}(t) \frac{C_{r_i}(t)x_{r_i}(t)}{C_{r_i}(t)x_{r_i}(t) + 1} + \mu_s(t - \tau_{si}) \frac{C_s(t - \tau_{si})x_{s_i}(t - \tau_{si})}{C_s(t - \tau_{si})x_s(t - \tau_{si}) + 1}, \quad (17)$$

where x_{s_i} represents the number of aircraft in the ramp queue that are bound for the i^{th} departure runway, τ_{gs} is the average unimpeded travel time from the gate to the spot, and τ_{si} represents the unimpeded travel time from the spot to the i^{th} runway. All departures are assumed to traverse through a single congested spot to reach the runway. While this assumption is found to be reasonable for CLT-NF, the approach can be easily extended to scenarios with multiple ramp queues.

The model parameters were determined from airport operational data that include flight tracks (Federal Aviation Administration 2018a), the actual pushback, in-air (wheels-off), landing (wheels-on)

and in-gate times, and gate assignments (OAG 2016), and meteorological conditions at the airport (Federal Aviation Administration 2018b). The data was obtained for a three month period (May-July 2016) for DFW and EWR, and a five month period for CLT (May-July 2015; May-Jun 2016). The dataset includes the periods of peak demand experienced in summer.

The service time distributions were obtained by computing the difference between successive exit times from a queue when there is pressure on the server. These distributions differ for different airports because of different procedures and fleet mix. The service time distribution for a departure runway server is conditioned on the number of landings and the weather at the airport (IMC/VMC). For example, the mean service rate μ (in aircraft per minute) for the the departure runway server (36R) at CLT under different weather conditions is estimated to be:

$$\mu_{d,r}(t)|_{VMC} = -0.11n_{a,r}(t) + 0.79; \text{ and } \mu_{d,r}(t)|_{IMC} = -0.11n_{a,r}(t) + 0.77, \quad (18)$$

where $n_{a,r}(t)$ denotes the number of landings on runway r in a 5-min time window containing t . The above expression for the mean service rate was empirically determined using linear regression with operational data (Badrinath et al. 2018). As one would expect, the mean service rate decreases as the number of landings on the runway increases (negative slope), and the intercept of the mean service rate is lower in IMC than during VMC. The service rate of the taxi-out ramp server is modeled as a function of the queue length of the taxi-in ramp queue, and vice versa. For example, the relationship between the taxi-out ramp service rates ($\mu_{d,s}(t)$) and the taxi-in ramp queue ($x_{a,s}$) at CLT is found to be:

$$\mu_{d,s}(t) = -0.033x_{a,s}(t) + 1.7. \quad (19)$$

The wait times of aircraft entering the queue are determined using the predictions of queue length and time-varying mean service rates (Badrinath, Li, and Balakrishnan 2018). Let $x(t_{in})$ be the predicted queue length at time $t = t_{in}$, $\mu(t)$ the mean service rate of the server, and Δt an appropriately small time-step. Then, the wait time (W) for an aircraft that enters the queue at t_{in} is estimated as follows:

```

Q = x(tin);
q = Q; t = tin; W = 0;
while q > 0 do
    | q = q - μ(t)Δt;
    | W = W + Δt;
    | t = t + Δt;
end
W = W + q/μ(t);
    
```

Algorithm 1: Calculation of wait time in a queue.

The predicted taxi time is the sum of unimpeded travel time and wait times in the different queues. To summarize, the proposed queuing model can predict surface queue lengths and taxi-out

times, given the pushback times, landing times and weather conditions, as well as the appropriate initial conditions on the queue lengths.

2.4. Predictive performance of the queuing model

Figure 7 shows a comparison of the predicted and observed departure runway queue lengths at DFW for a typical day. The data corresponds to a time-based definition of queue length, in which an aircraft is said to be in the runway queue if it has spent unimpeded gate-to-runway time after pushback but is yet to take-off. Our analysis has found that this time-based definition is consistent with the observed physical queue. The taxi-out times for this particular day, averaged over 15-min windows, are shown in Figure 8. These figures show a good match between the predictions and observed values. We also note that errors in the taxi-out time predictions appear to be correlated with prediction errors in queue lengths.




Figure 7 Comparison of model predictions and observed queue lengths on a typical good weather day at DFW (07/20/2016).

Table 2 shows the aggregate error statistics of taxi-out time prediction for individual flights, computed for an independent test set. Flights with taxi-out times greater than 50 min were not




Figure 8 Taxi-out times averaged over 15-min intervals on 07/20/2016.

included while computing the statistics. The mean errors (ME) are found to be small relative to the mean taxi-out times and negative at all three airports. The mean absolute error (MAE) of the taxi-out time is similar to the results obtained from other probabilistic queuing models (Simaiakis and Balakrishnan 2015). A high value of MAE indicates that the takeoff sequence does not necessarily follow the first-come-first-served assumed by the queuing model. Figure 9 compares the predicted and actual taxi-out time distributions. We note that the right tail of the taxi-out time distribution is not captured by the queuing model, resulting in a negative value of the mean error. This right-tail is more pronounced for EWR and CLT, where in addition to queuing delays, surface gridlock can cause further delays that are not accounted for by the queuing model. Table 3 shows the prediction error statistics for individual flights grouped by different quartiles of the actual taxi-out time. For flights in the first quartile (0-25th percentile), the MAE is small and ME is positive. The model slightly over-predicts taxi-out time in this range because it assumes that the unimpeded time is the 10th percentile of the taxi-out time distribution. For the flights belonging to the upper quartile (75th-100th percentile), the magnitude of the ME is negative and large. From the point of view of departure metering, positive errors correspond to over-prediction of the taxi-out time, leading to overly aggressive hold decisions, loss in runway utilization, and delays in wheels-off time. By contrast, negative taxi-out time errors will yield lower metering benefits, but there would be no loss in throughput and no additional delays.

Airport	Configuration	Number of departures	Taxi-out time (min)			Percentage of flights with $ \text{error} < 5$ min
			Mean	ME	MAE	
EWR	NF	2,596	21.3	-1.4	4.9	61.9
	SF	6,877	21.5	-1.9	5.3	61.1
DFW	NF	2,769	18.5	-1.2	3.9	72.3
	SF	19,207	17.4	-1.1	4.2	71.5
CLT	NF	7,464	20.1	-1.4	4.4	69.0
	SF	6,120	20.1	-1.9	5.2	61.0

Table 2 Error statistics of the taxi-out time predictions from the analytical queuing model on a test set.

Airport	Configuration	0-25 th %ile		25 th -50 th %ile		50 th -75 th %ile		75 th -100 th %ile	
		ME	MAE	ME	MAE	ME	MAE	ME	MAE
EWR	NF	2.9	3.6	0.8	3.5	-1.7	4.6	-7.2	8
	SF	2.5	3.2	0.7	3.4	-1.5	4.4	-9.2	10.1
DFW	NF	1.4	2.5	0.3	2.9	-0.7	3.5	-5.1	6.4
	SF	1.7	2.5	0.8	2.9	-0.4	3.5	-5.8	7.4
CLT	NF	1.8	2.9	0.3	3.2	-1.3	3.9	-6.3	7.5
	SF	2.5	3.3	0.5	3.5	-2.3	4.6	-7.7	9.1

Table 3 Error statistics for the taxi-out time predictions for individual flights from the queuing model based on the test set for different quartiles of the actual data.




Figure 9 Comparison of the predicted and actual taxi-out time distributions.

3. Departure metering

During periods in which demand exceeds capacity, queues build up on the airport surface, leading to an increase in taxi-out times. The idea behind departure metering is to tactically assign hold-time for departures at the gate during periods of peak traffic, so that smaller queues are formed on the airport surface, leading to a reduction in taxi-out time. The challenge is to assign appropriate hold-times to maximize benefits in terms of taxi-out time reduction, while maintaining runway utilization. Departure metering is an integral part of the FAA’s Terminal Flight Data Manager (TFDM), which is a NextGen initiative that is scheduled to be deployed from 2020 (Federal Aviation Administration 2018c). Several algorithms have been previously proposed to determine the appropriate hold times at the gate (Simaiakis et al. 2014, Pujet, Delcaire, and Feron 2000, Badrinath and Balakrishnan 2017, Khadilkar 2013). This paper considers two approaches to departure metering: The first one based on NASA’s ATD-2 logic, and the second based on an optimal control formulation.

3.1. ATD-2 logic for departure metering




Figure 10 A timeline diagram of the ATD-2 logic for departure metering. UTOT indicates undelayed takeoff time.

The ATD-2 logic for departure metering computes a gate-hold time for each flight based on its predicted taxi-out time as follows (Verma et al. 2018):

$$TOBT = \max(EOBT, TTOT - UTT - Y), \quad (20)$$

where $TOBT$ is the Target-Off-Block-Time or the new gate release time decided by the controllers, $TTOT$ is the Target Take-Off-Time, UTT is the unimpeded time to take-off that depends on the

gate-runway pair, and Y is the excess queue time buffer (Verma et al. 2018). The target take-off time is computed by adding the predicted taxi-out time to the flight's EOBT. In this paper, we use the queuing model presented earlier to obtain the taxi-out time prediction for each flight. In other words, the hold time assigned to each flight is the predicted wait time in queue for that flight minus the excess queue time buffer. The idea is to transfer the predicted wait time in the queues to a gate-hold time, saving fuel. The excess queue time buffer (Y) is subtracted to account for errors in the taxi-out time prediction. The buffer value (Y) is a constant design parameter that is chosen for each airport and runway configuration. It is important to choose the optimal buffer parameter. If the buffer parameter is too high, it will lead to decreased benefits; if the buffer parameter is too low, it could lead to reduced runway utilization. A timeline diagram of the ATD-2 logic is shown in Figure 10.

Given the weather condition (IMC/VMC), landing times for arrivals, runway assignments and EOBT time for departures, the queuing model can predict the taxi-out time for each flight based on the current queue length at an airport. The hold time is then assigned to each flight based on its predicted taxi-out time. To have higher predictability in the system, the hold decisions are made T_p minutes prior to a flight's EOBT (where T_p is the planning horizon).

3.2. Simulation environment for the evaluation of departure metering algorithms

The departure metering approaches are evaluated using a simulation of airport surface operations. The simulators are based on discrete queuing network models (as described in Section 2), with the service time for each server being sampled from an empirical distribution. The empirical service time distributions are a function of the airport weather and traffic (such as number of landings on the runway and traffic on the ramp) as we had discussed earlier. The simulations are repeated multiple times to obtain consistent statistics (a Monte Carlo simulation with 20 runs). We consider a 15 day period for each airport-runway configuration pair to evaluate the benefits of departure metering. Table 4 shows the error statistics for the taxi-out time predictions from the simulator in the baseline case without any metering by comparing it with the actual data. Although we consider a 15-day period for each airport-runway-configuration pair, the total number of departures vary between NF and SF, because the total time duration operating in a particular configuration is different. The results indicate that the errors are small relative to the mean taxi-out times, but their magnitudes are slightly larger than the analytical queuing model.

3.3. Departure metering benefits with accurate EOBTs

In this section, we assume that we have accurate EOBT information, that is, the estimate of the push-ready time published by the airlines is accurate. This is a reasonable assumption for scenarios in which the hold time is assigned to flights when the pilot calls ready (Simaiakis et al. 2014, Verma

Airport	Configuration	Number of departures	Taxi-out time (min)			Percentage of flights with $ \text{error} < 5$ min
			Mean	ME	MAE	
EWR	NF	4,117	21.6	0.4	5.4	60.4
	SF	4,493	21.0	0.5	5.4	59.0
DFW	NF	6,377	18.7	-0.4	4.2	70.6
	SF	9,212	17.3	0.6	4.9	64.4
CLT	NF	6,447	20.1	1.1	4.6	64.2
	SF	4,493	19.7	-1.0	5.0	63.8

Table 4 Error statistics for the baseline taxi-out time predictions from the queuing simulation over a 15 day period.

et al. 2018), including the current phase of the ATD-2 demonstration. However, in the long-term, the plan is to assign gate holds several minutes (say 10-30 min) prior to the flight's EOBT. At time T_p prior to a flight's EOBT, the taxi-out time for a flight is predicted from the analytical queuing model using the EOBT information for all the departures, landing times for future arrivals, weather condition and current traffic at the airport. For the purpose of evaluating the departure metering benefits with accurate EOBTs, we consider the flight's EOBT to be equal its Actual Off-Block Time (AOBT) obtained from historical data. The gate release time (TOBT) for the flight is then determined using Eq. (20). The flights are released at TOBT in the discrete queuing simulation of the airport. In the field demonstrations at CLT, NASA uses a different model for taxi-out time prediction, a microscopic node-link model of the airport, instead of a queuing model (Verma et al. 2018).





Figure 11 Metering simulation results for different excess queue time buffer values. The plots show averaged values over a 15 min time window for a typical day at DFW in North-Flow configuration (07/20/2016).

Figure 11 shows the simulation results of departure metering with a planning horizon of 20 min for different values of the excess queue time buffer values (Y), on a typical day at DFW. The reduction in taxi-out time is correlated with the hold-time, which in turn is larger during periods when the baseline taxi-out time is high. As the excess queue time buffer increases, the average hold-time decreases, leading to lower taxi-out time reduction benefits. On the other hand, a lower excess queue time buffer might lead to aircraft being held longer than necessary, leading to a delayed



(a) Average taxi-out reduction and hold time for different buffers (min).

(b) Hold time distribution for the flights that were held.

Figure 12 Departure metering statistics for DFW-SF. The statistics were calculated over 9,212 flights in a 15 day period.

wheels-off time (Figure 11(c)). The excess queue time buffer is therefore an important parameter that influences the taxi-out time reduction and airport throughput. In the field demonstrations at CLT conducted by NASA, ramp managers are allowed to pick the excess queue time buffer based on their experience (Verma et al. 2018). We present a methodology for systematically selecting the excess queue time buffer parameter.

The wheels-off delay for a flight because of departure metering is given by,

$$\text{wheels-off delay} = (\text{wheels-off time})_{\text{metering}} - (\text{wheels-off time})_{\text{baseline}} \quad (21)$$

$$= \left((\text{taxi-out time})_{\text{metering}} + \text{hold-time} + \text{AOBT}_{\text{baseline}} \right) \quad (22)$$

$$- \left((\text{taxi-out time})_{\text{baseline}} + \text{AOBT}_{\text{baseline}} \right) \quad (23)$$

$$= \text{hold-time} - \left((\text{taxi-out time})_{\text{baseline}} - (\text{taxi-out time})_{\text{metering}} \right) \quad (24)$$

$$= \text{hold-time} - (\text{taxi-out time reduction}) \quad (24)$$

The difference between the hold-time and taxi-out time reduction is equal to the wheels-off delay. A positive value of this change indicates that the runway is underutilized. Figure 12(a) shows the average hold time and taxi-out time reduction for different excess queue time buffers in DFW-SF. As the buffer value increases, the hold-time as well as reduction in the taxi-out time decreases. We can see that at small values of the excess queue time buffer, the wheels-off delay is positive, indicating that the runway is underutilized. The optimal excess queue time buffer is chosen to achieve as large a taxi-out time reduction as possible while not losing runway utilization. For the 20 min horizon considered here, the optimal excess queue time buffer is 5 min for DFW-SF, with an average taxi-out time reduction of 2.1 min.

The optimal buffer size and taxi-out time reduction vary by airport and runway configuration. Table 5 shows the optimal buffer size and taxi-out time reduction from the simulations of departure

Parameter	EWR		DFW		CLT	
	NF	SF	NF	SF	NF	SF
Optimal buffer (min)	7	8	5	5	7	5
Mean taxi-out time reduction(min)	2.0	1.7	2.0	2.0	2.7	1.9
Mean hold-time (min)	2.1	1.7	2.0	2.1	2.8	1.9
Mean wheels-off delay (min)	0.07	0.02	0.05	0.02	0.10	0.06
Percentage of flights held	61%	56%	58%	59%	64%	53%
Mean hold time of the flights held (min)	3.4	3.1	3.5	3.5	4.3	3.6

Table 5 Effects of departure metering with 20-min planning horizon without EOBT uncertainty.

metering, assuming perfect EOBT information. As one would expect, the optimal buffer size is higher for airports with higher taxi-out time prediction errors. Although the taxi-out delays at EWR and CLT are larger than those at DFW, the reduction in taxi-out time is about the same due to differences in prediction errors. At all the airports, about 60% of flights are held at their gates, and the mean hold time of these flights is about 3 min. The distribution of the hold time for the flights that were held is shown in Figure 12(b) for DFW-SF.

3.3.1. Gate conflicts A gate conflict occurs when an arriving flight needs to wait for the gate occupied by another aircraft that is supposed to depart but which is currently being held. Gate conflicts could potentially decrease the benefits of departure metering, since an arrival may need to use a gate at which a departure is being held. Figure 13 shows the cumulative distribution of the time between pushback and the next arrival time at the gate for the current operations at DFW. It shows that only 20% of the departing flights have an arriving aircraft that uses the same gate within 20 min of the departing flight's pushback time. Since the simulations indicate that the average hold time for the flights held at the gate is 3.5 min and maximum hold time is 14 min (corresponding to optimal buffer size), we do not expect a significant reduction in the taxi-out time savings due to gate conflicts.




Figure 13 Cumulative distribution of the time between the actual pushback and next arrival at a gate.

4. Impact of EOBT uncertainty

The results for the taxi-out time prediction and departure metering presented so far assumed that the EOBT information provided by the airlines was accurate. However, the EOBTs published by airlines often deviate from the actual push-ready times, motivating the study of the impact of EOBT errors on departure metering.

4.1. Evaluating EOBT uncertainty

EOBT information was obtained from the FAA’s Traffic Flow Management System (TFMS) for EWR, DFW and CLT. Based on the data availability, we considered the following periods for our analysis: Jan 01 - Nov 11, 2018 (EWR), May 01 - July 31, 2018 (DFW) and Dec 13, 2017 - Feb 12, 2018 (CLT). Flights departing between 9AM and 12PM (local time) were removed from CLT’s dataset to eliminate any effects of the ATD-2 departure metering demonstration.




Figure 14 Message time of the EOBT information relative to the EOBT time (Airline A).

Figure 14 shows the distribution of the time when the EOBT information was published for a flight relative to its EOBT for a major airline (*Airline A*) at the three airports. Most of the EOBT messages appear to be published within 40 min prior to the flight’s EOBT (median: 20-23 min), with an additional spike at around 73 min. Further analysis indicates that airlines publish the Initial Gate Time of Departure (IGTD) or the original scheduled gate pushback time prior to the 40 min mark (Badrinath et al. 2018). Table 6 shows the EOBT update statistics for the three airports.

Parameter	Mean			Median			σ		
	EWR	DFW	CLT	EWR	DFW	CLT	EWR	DFW	CLT
Message time – EOBT (min)	-22.7	-22.2	-18.6	-23.0	-23	-20.1	5.5	5.5	8.3
Time between updates (min)	2.7	4.4	2.8	1.6	1.8	2.0	3.3	5.3	2.6
No. of EOBT messages per flight	10.3	9.8	13.3	9	8	10	5.3	5.8	10.2

Table 6 EOBT update statistics computed for lookahead horizons lesser than 40 min (Airline A).

The EOBT information is also dynamically updated by airlines at irregular intervals until the flight pushes back from the gate (FAA Surface CDM Team 2012). The EOBT updates are a concern




Figure 15 Time between EOBt updates for lookahead time lesser than 40 min (Airline A).




Figure 16 Distribution of number of EOBt updates received by a flight which had its EOBt published (Airline A).

for the decision makers who need a ‘stable’ information for making decisions. Figure 15 shows the distribution of the time between EOBt updates for a flight, considering messages within 40 min of the flight’s EOBt. Only 28% of the flights operating out of EWR currently have their EOBt published. However, all flights with a planned departure are expected to publish EOBt information in the near future (FAA Surface CDM Team 2012). Figure 16 shows the distribution of the number of EOBt updates for a flight, given that EOBt information for it was published at least once. The median number of updates is between 8 and 10 depending on the airport.

The EOBt error was calculated by subtracting the ASPM-derived Actual Off Block Time (AOBT) for that flight from the published EOBt (Federal Aviation Administration 2018b). These errors were compared over different lookahead time horizons (the difference between the EOBt message transmission time and the EOBt value published at that time). Figure 17 shows the results of fitting a Probability Density Function (PDF) to the EOBt error distributions for different lookahead times, obtained for Airline A at EWR, DFW and CLT. The key error statistics are presented in Table 7. As expected, the standard deviations of the distributions increase with increasing lookahead time. The EOBt error distributions also vary by airport. It is worth noting that the standard deviation of the EOBt error, even for lookahead times of 10-20 min, is of the same order of magnitude as the taxi-out time prediction errors from the queuing model. It is therefore important to consider the




Figure 17 EOBT error distribution for different lookahead horizons (Airline A).

uncertainty in EOBT information while evaluating the benefits of departure metering. The EOBT error distributions for all lookahead horizons are negatively skewed. The distributions for a 20-min horizon show that flights generally do not pushback more than 10 min prior to their EOBT; however, there are flights that get delayed beyond 20 mins of the published EOBT.

The EOBT error distributions vary by airline, since each airline does its EOBT computation internally. Table 8 shows the EOBT error statistics for two other airlines at EWR (*Airline B* and *Airline C*). The errors are significantly higher for these two airlines compared to Airline A. Figure 18 presents the EOBT update distributions for Airline B at EWR, and shows that Airline B tends to update its EOBT information less frequently than Airline A.

Lookahead time (min)	Mean error (min)			Median error (min)			σ error (min)		
	EWR	DFW	CLT	EWR	DFW	CLT	EWR	DFW	CLT
10	1.0	-1.8	1.5	1.1	-0.4	1.8	3.1	5.9	3.9
20	1.5	-1.8	1.4	1.8	-0.2	1.7	3.5	6.1	3.9
30	0.0	-6.2	-0.1	2.0	-2.0	2.2	8.2	13.7	8.9
40	-1.5	-12.9	-4.5	0.0	-7.0	0.0	8.9	18.4	14.7

Table 7 EOBT error statistics for different lookahead times (Airline A).

Lookahead time (min)	Mean error (min)		Median error (min)		σ error (min)	
	Airline B	Airline C	Airline B	Airline C	Airline B	Airline C
10	-3.9	-3.9	-1.0	0.2	11.4	9.8
20	-7.7	-5.1	-5.0	0.3	14.3	10.9
30	-7.9	-7.2	-5.0	0.5	14.8	9.8
40	-9.1	-7.5	-4.0	0.7	17.0	13.2

Table 8 EOBT error statistics for Airline B and Airline C at EWR computed for different lookahead horizons.




Figure 18 EOBT information update statistics for Airline B at EWR.

4.2. Parametric analysis of the impacts of EOBT uncertainty on departure metering

The uncertainty in EOBT information impacts departure metering for two reasons:

- *Reduced accuracy of taxi-out time predictions*: Departure metering algorithms require an estimate of taxi-out times to determine gate-hold times. The taxi-out time predictions need the EOBTs as inputs (for example, to a queuing model). Errors in EOBT result in taxi-out time prediction errors, leading to inefficient gate-hold times.

- *Noncompliance with assigned gate release times*: The pushback time recommendation (TOBT) is given by the EOBT time plus the gate-hold time. However, a flight may not be ready to push at its TOBT due to errors in its EOBT, resulting in a loss in runway throughput.

We conduct a parametric study of the impacts of inaccurate EOBTs on departure metering using the ATD-2 logic. We assume that the EOBT error distributions are Gaussian with a zero mean, and we vary the standard deviations. This methodology allows us to investigate the impact of different error distributions in a systematic way. The range of the parametric space for the standard deviation of the error distribution follows from the operational data. The EOBT time is synthetically computed for every flight in the data-set by adding a random variable sampled from the EOBT error distribution to the AOBT (in other words, $EOBT = \text{error} + AOBT$).

For a flight with an uncertain EOBT, we can have one of the following three cases:

1. $Push\text{-ready}\ time < EOBT$: In this case, the metering logic would recommend a gate-hold until the EOBT time even if the predicted excess queue time is equal to zero. While this strategy may improve predictability, flights would have to hold unnecessarily at the gate, resulting in a loss in runway throughput. To prevent unnecessary holds, the metering logic is modified to let flights pushback at the push-ready time if the predicted excess queue time (along with the buffer) is equal to zero.

2. $EOBT \leq push\text{-ready}\ time \leq TOBT$: In this case, the flight waits and then pushback occurs at its TOBT. The gate-hold time would be expected to translate into a reduction in taxi-out time.

3. $TOBT < \text{push-ready time}$: In this case, the flight is not ready to pushback at its TOBT and the pushback occurs later. Once again, there is a loss of predictability and a potential loss in runway throughput because the flight did not pushback when expected.

The optimal buffer size in the ATD-2 logic can be adjusted to account for the EOBT uncertainties, in addition to the taxi-out time prediction errors. The optimal excess queue buffer is the smallest value of the queue buffer for which there is no increase in the wheels-off time. Figure 19(a) shows the optimal excess queue buffer time for different EOBT distributions and two different planning horizons (10-min and 20-min) for DFW-SF. For a particular planning horizon, the optimal excess queue buffer increases with an increase in EOBT error and then saturates beyond a particular point. This behavior is as expected, since the taxi-out time prediction error increases with EOBT uncertainty and needs to be compensated with a larger excess queue time buffer. For higher EOBT errors, a larger buffer size is required at the 20-min planning horizon than the 10-min horizon, due to the expected decrease in prediction performance.

There is an inherent trade-off in choosing the planning horizon: A larger planning horizon is preferable since it leads to more predictability, but it yields lower departure metering benefits. One would ideally identify the optimal value of the planning horizon by evaluating the tradeoff between predictability and fuel savings; however, it is difficult to do so since the value of predictability is not known. In this paper, we consider the two values of planning horizon (10 and 20 min) that have been proposed for ATD-2.

We can also evaluate the departure metering benefits for different EOBT error distributions using the optimal excess queue buffer time. The taxi-out time savings decreases as the standard deviation of the EOBT error increases (Figure 19(b)). At a 20-min planning horizon, with the current level of EOBT uncertainty at DFW (standard deviation of 6 min), there is a 50% decrease in taxi-out time reduction benefits compared to the case with perfect EOBT information.

EOBT error σ (min)	EWR		DFW		CLT	
	NF	SF	NF	SF	NF	SF
0	2.0 (5)	1.7 (8)	2.0 (5)	2.1 (5)	2.7 (7)	1.9 (5)
2	1.6 (6)	1.3 (8)	1.7 (6)	1.8 (6)	2.3 (8)	1.6 (5)
4	1.2 (8)	1.2 (9)	1.2 (8)	1.3 (8)	1.8 (10)	1.4 (6)
6	0.9 (10)	1.2 (10)	0.9 (10)	1.0 (10)	1.3 (12)	1.3 (7)
8	0.7 (11)	0.9 (12)	0.7 (11)	0.8 (11)	1.1 (13)	1.1 (8)
10	0.5 (12)	0.9 (13)	0.5 (12)	0.6 (12)	1.0 (13)	1.1 (9)
12	0.5 (12)	0.9 (13)	0.5 (12)	0.5 (12)	0.9 (13)	0.9 (10)

Table 9 Average taxi-out time reduction in minutes for different EOBT error distributions with a planning horizon of 20-min. The optimal buffer for the ATD-2 logic is indicated in brackets.

Table 9 shows the average taxi-out time reduction and optimal buffer for different levels of EOBT uncertainty at different airports and a 20-min planning horizon. The trends in the reduction of

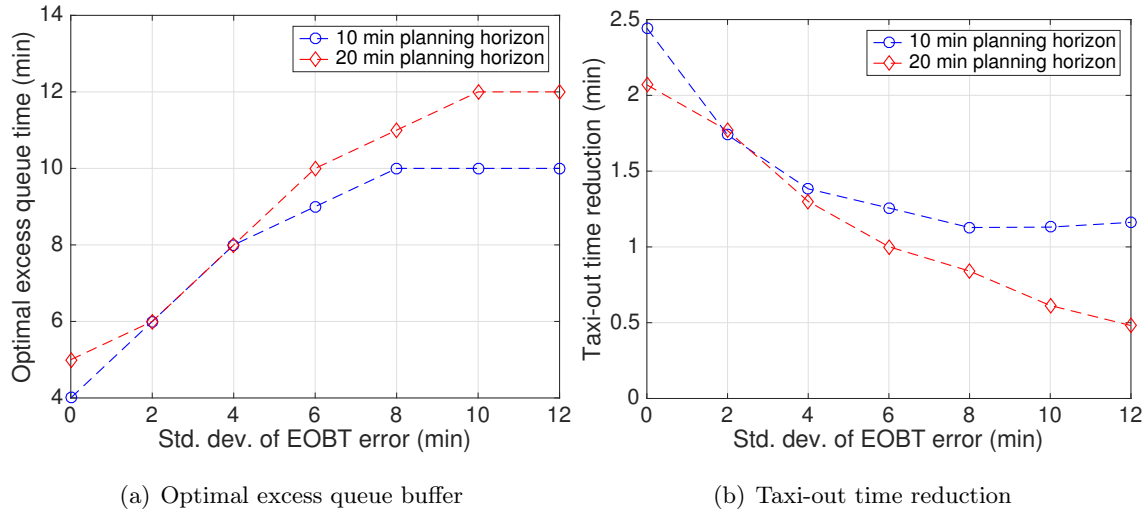


Figure 19 Optimal excess queue buffer and taxi-out time reduction for DFW-SF.




Figure 20 Annualized taxi-out time reduction and fuel burn reduction per flight using the ATD-2 logic for different levels of EOBT uncertainty.

metering benefits with EOBT uncertainty are consistent across all the airports. Figure 20 shows the annualized average taxi-out time and fuel burn reductions as a function of the EOBT error. These calculations account for the relative frequency of different runway configurations and the average fleet mixes at the different airports. The values of the average taxi fuel flow rate per flight were found to be 0.20 kg/s for EWR and DFW, and 0.17 kg/s at CLT (due to the higher fraction of regional jets) (ICAO 2018). Table 10 shows the estimated annualized total fuel burn saving at the three airports obtained by multiplying the average fuel burn savings per flight with the total number of operations in 2017.

4.3. An optimal control approach for departure metering

We propose an approach based on optimal control technique to departure metering in the presence of EOBT uncertainty, by leveraging the analytical queuing models developed in Section 2. The objective of the formulation is to regulate the rate of aircraft pushbacks so as to achieve smaller queue lengths on the airport surface, while maintaining runway throughput.

Airport	EOBT error σ (min)						
	0	2	4	6	8	10	12
EWR	4.9	4	3.4	3.2	2.7	2.6	2.5
DFW	8.4	7.2	5.3	4	3.3	2.4	2
CLT	6.6	5.5	4.6	3.8	3.2	3.1	2.7

Table 10 Annualized fuel burn savings ($\times 10^6$ kg) using the ATD-2 departure metering logic.

Let $\mathbf{x}(t) \in \mathbb{R}^p$ be a vector of taxi-out queue lengths on the airport surface at any time instant t . Let $\mathbf{d}(t) \in \mathbb{R}^q$ be the departure demand rate, with the elements representing the demand rate to each runway. The departure demand corresponds to the push-ready time for departures, and is obtained from the EOBT information provided by the airlines. These rates are obtained over 5-min windows. Let $\mathbf{u}_d(t)$ be the pushback rate at the gate that is assigned by the controllers, and $\mathbf{h}(t)$ be the number of aircraft held at the gate due to departure metering. The number of holds at any instant, $\mathbf{h}(t)$, is the difference between the total demand and total pushbacks until time t . The dynamics for the number of holds is then given by:

$$\mathbf{h}(t) = \int_0^t \mathbf{d}(x)dx - \int_0^t \mathbf{u}_d(x)dx \quad (25)$$

$$\dot{\mathbf{h}}(t) = \mathbf{d}(t) - \mathbf{u}_d(t) \quad (26)$$

The control objective is to minimize the length of taxi-out queues while maintaining throughput, the state variables are the taxi-out queue lengths on the airport surface and the number of holds, and the control variable is the departure pushback rate ($\mathbf{u}_d(t)$). The problem formulation for an airport queuing network is then as follows:

$$\min_{\mathbf{u}_d(t)} \int_0^T (\mathbf{x}^T Q \mathbf{x} + \mathbf{h}^T R \mathbf{h}) dt \quad (27)$$

$$\text{Subject to:} \quad (28)$$

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}(t), \mathbf{x}(t - \tau_1), \dots, \mathbf{x}(t - \tau_m), \mathbf{u}_d(t - \tau_{m+1}), \dots, \mathbf{u}_d(t - \tau_w), t) \quad (29)$$

$$\dot{\mathbf{h}} = \mathbf{d}(t) - \mathbf{u}_d(t) \quad (30)$$

$$0 \leq x_i, h_i; \quad 0 \leq u_{di} \leq u_m; \quad i = 1, 2, 3, \dots, w \quad (31)$$

$$u_{di}(t) = g_i(t), \quad t \in [-\tau_{di}, 0]; \quad i = 1, 2, 3, \dots, w \quad (32)$$

$$x_i(t) = \phi_i(t), \quad t \in [-\tau_{ki}, 0], \quad \mathbf{h}(0) = \mathbf{h}_0; \quad i = 1, 2, 3, \dots, w \quad (33)$$

Here, T is the optimization window over which the cost needs to be minimized, and $Q \in \mathbb{R}^{p \times p}$ and $R \in \mathbb{R}^{q \times q}$ are constant weighting matrices. Equations (29) and (30) specify the dynamics for the length of the queues and number of holds. The inequalities in (31) impose nonnegativity constraints on the number of holds, queue lengths and pushback rate. The delay differential equations also require an initial history, specified by Eqs. (32)-(33). The cost function (27) penalizes a weighted

sum of the square of the queue lengths and number of holds. The queue length is penalized to reduce taxi-out times, while the holds are penalized to avoid having large holds and to maintain runway throughput.

In Section 2, we saw that the dynamics of surface queues depend on the service time of several servers. The service time distributions for the departure runway queues depend on the landing rates and meteorological conditions, which are assumed to be known. The service time distribution of the taxi-out ramp queue requires the number of taxi-in flights on the ramp (in the case of CLT-NF). The taxi-in ramp queue length is pre-computed using the queuing model with the EOBTs and arrival times, and is used to determine the service rate for the taxi-out ramp queue. This service rate is used for the taxi-out ramp queue dynamics in the optimal control problem. In this manner, the interdependencies between departures and arrivals in the ramp are simplified so that only the arrivals impact the departures. This methodology reduces the number of constraints in the optimal control problem. To speed up the computation further, the optimal pushback rate is computed only when the runway queue length at T_p min in the future is predicted to be greater than 2.

The optimal control problem is solved in a receding horizon framework, accounting for the current state of the airport surface. The entire day is divided into 5-min intervals. At the beginning of each interval, pushback rate decisions are made for $[T_p, T_p + 5]$ min, where T_p is the planning horizon. The optimal control problem is solved for $[T_p, T_p + T]$ where T is the optimization window over which the cost is minimized, and the pushback rate is implemented for the first 5 min. The initial conditions for the optimal control problem at T_p min into the future are obtained using the queuing network model with the current state as the input. The number of aircraft that can be released during each 5 min window (n) is determined from the pushback rate. The first n aircraft in the 5 min window are released as per the optimal control decision and rest of the aircraft are pushed to the beginning of the next time window, awaiting decision for release. In contrast to the ATD-2 logic which assigns and freezes the hold time for flights that have an EOBT T_p min ahead, the optimal control approach only specifies the flights that need to be released in $[T_p, T_p + 5]$ time window, and pushes the remaining flights to the next time window. Consequently, aircraft can be pushed multiple times to the next time window but the release decision is made T_p minutes ahead.

The optimal control problem is solved numerically rather than analytically due to the challenges posed by time delays and nonlinear dynamics. The optimal control problem is solved by discretizing the state and control variables. The equations governing the dynamics are discretized using a first-order Euler method. The discretized control problem is then transformed into a non-linear programming problem (NLP) that can be solved using a standard solver in MATLAB.

We demonstrate the benefits of the optimal control approach by simulating 3 days of operations in DFW-SF and CLT-NF. The queuing models of these two airport-configuration pairs differ significantly in complexity. The DFW dataset consisted of 2,708 departures and CLT dataset consisted

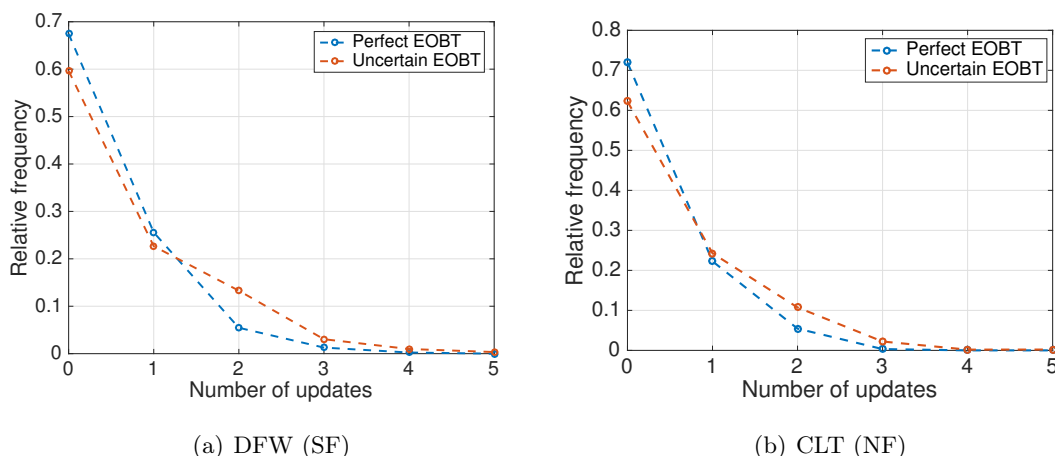


Figure 21 Relative frequency of the number of updates for the release time for the flights that were held at the gate using the optimal control policy.

of 1,817 departures. We penalize the departure runway queue in the cost function with a weight equal to one, and impose a weight of 0.4 for the holding term ($Q = \mathbb{I}_p, R = 0.4\mathbb{I}_q$). The weight of the holding term was determined in a similar way as we had determined the optimal queue buffer for ATD-2 logic, to ensure that the mean wheels off delay is close to zero. We consider T , the optimization window over which the cost is minimized to be 30-min for CLT and 60-min for DFW. The longer horizon yielded better departure metering benefits at DFW, because of the longer departure banks.

Table 11 compares the benefits of the optimal control approach and the ATD-2 logic. While the ATD-2 logic yields more benefit with perfect EOBT information, the optimal control approach performs better in the presence of uncertainty. The benefits in terms of taxi-out time reduction with the optimal control approach in the presence of EOBT uncertainty are about 66% higher for DFW-SF and 33% higher for CLT-NF. Figure 21 shows the relative counts of the number of times a flight is pushed to the next 5 min interval, given that it was asked to hold at least once. Most of the flights that are held at the gate receive at most two updates.

Parameter	DFW-SF				CLT-NF			
	Perfect EOBT		Uncert. EOBT		Perfect EOBT		Uncert. EOBT	
	Opt.	ATD-2	Opt.	ATD-2	Opt.	ATD-2	Opt.	ATD-2
Mean taxi-out time reduction(min)	1.4	1.9	1.0	0.6	1.3	2.6	1.2	0.9
Mean hold-time (min)	1.5	1.9	1.1	0.7	1.5	2.8	1.3	1.0
Mean wheels-off delay (min)	0.1	0.1	0.1	0.1	0.2	0.1	0.1	0.0
Fraction of flights held	0.3	0.6	0.2	0.1	0.3	0.6	0.2	0.1
Mean hold time of flights held (min)	4.9	3.5	5.3	6.3	4.5	4.3	6.2	6.8

Table 11 Comparison of the departure metering benefits with and without EOBT uncertainty (assuming $\sigma = 6$ min).
 Opt. refers to results from the optimal controller.

Table 12 compares the computational times of the optimal control approach and the ATD-2 logic, based on the simulation of a single day of airport operations. The computations were done on a computer with 2.8 GHz Intel Core i7 processor with 16 GB memory, and the scripts were run on MATLAB. The computational time of the ATD-2 logic corresponds to deciding the gate-holds for one or more flights that have the same EOBT. The computational time can vary slightly depending on the number of flights that share a particular EOBT. The computational time indicated for the optimal control approach is the time required to make a decision for every 5-min window in the receding horizon framework. The time duration to solve the optimal control problem can vary based on the instance of the optimization problem. In general, the ATD-2 logic is orders of magnitude faster than the optimal control approach in terms of computational time. However, the optimal control approach is still practical: The mean time for planning a 5-min window is around 2 sec, and maximum computation time is around 40 sec. Finally, the computational time for CLT is higher than that for DFW, because of the more complex queuing network.

Airport	Mean comp. time (sec)		Max. comp. time (sec)	
	ATD-2	Opt. Ctrl.	ATD-2	Opt. Ctrl.
DFW-SF	0.5×10^{-3}	1	17×10^{-3}	31
CLT-NF	0.9×10^{-3}	3	24×10^{-3}	39

Table 12 Computational times for the ATD-2 logic and the optimal control approach.

5. Discussion and extensions

We discuss some of the open questions as well as extensions of this work.

5.1. Integrated control of arrival and departure queues

The optimal control formulation above only considered the taxi-out queues on the surface. Since we have an integrated model for taxi-out and taxi-in flights (as in CLT-NF), it would be a straightforward extension to account for taxi-in queues as well. The objective of the integrated arrival-departure control problem would be to minimize the weighted sum of all queue lengths on the airport surface by controlling the pushback rate and the arrival rate, accounting for the number of holds in the air (arrivals) and at the gate (departures). Additionally, the weights in the cost function could be chosen to penalize airborne holds more than the ground holds. In addition to assigning gate-holds, the optimal control approach can be extended to assign runways. However, there is less flexibility in assigning runways since there are constraints arising from the runway length required by aircraft, departure fix assignments, etc.

5.2. Robust control techniques

One could explore concepts from robust control theory to determine an algorithm that explicitly accounts for model uncertainties with theoretical guarantees. The dynamics of the queuing network being represented as a differential equation makes it amenable to use standard techniques in robust control such as robust model predictive control, sliding mode controllers, or \mathcal{H}_∞ controllers. The underlying challenges of using these approaches are that the nonlinearities in the queuing dynamics and the time delay.

The other source of uncertainty is the error in the predicted arrival rate and weather conditions. For airports that have runways under mixed operations (such as CLT-NF), we have assumed a knowledge of the landing times of arrivals in order to estimate the runway capacity for the departures. It is worth exploring the impact of arrival rate uncertainty as well.

5.3. Fairness of departure metering

The parameters of the departure metering algorithms were tuned to ensure that there was no loss in throughput, that is, no net delay for the departures from the airport. However, depending on the hold-time assigned to the flights, the take-off order in the presence of metering might be different compared to the baseline. This could result in one flight experiencing less delay at the expense of more delay for another flight.

Table 13 shows the mean absolute displacement in the takeoff order with metering. Here, “displacement” refers to the rank difference of the takeoff order at each runway with departure metering, compared to the baseline case without departure metering. The buffer parameter ensures that there is no loss in throughput, and that the average takeoff delay is zero. However, a higher value of mean absolute displacement indicates that some flights take off earlier in the metering case compared to the baseline, while some others take off later. Ideally, fairness suggests that the takeoff order should be the same in the metering case as in the baseline. Table 13 shows that the mean absolute displacement is higher in the presence of EOBT uncertainty.

Category	DFW-SF				CLT-NF			
	Perf. EOBT		Uncert. EOBT		Perf. EOBT		Uncert. EOBT	
	Opt.	ATD-2	Opt.	ATD-2	Opt.	ATD-2	Opt.	ATD-2
Mean absolute displacement of flights that are not held	0.2	0.0	0.5	0.6	0.4	0.2	0.6	0.8
Mean absolute displacement of flights that are held	1.5	1.2	1.9	4.0	1.6	1.7	2.0	3.9
Mean absolute displacement of all flights	0.6	0.7	0.8	1.0	0.8	1.2	0.9	1.3

Table 13 Mean absolute displacement of the takeoff order with and without EOBT uncertainty (assuming $\sigma = 6$ min). Opt. refers to results from the optimal control approach.

5.4. Incentives to improve EOBT accuracy

The analysis presented in this paper quantifies the benefits of more accurate EOBTs in terms of taxi-out time savings. The standard deviation of the EOBT error distribution is around 4 min at a 20-min time horizon for EWR and CLT, while it is 6 min at DFW (Table 7). If the airlines were to improve the EOBT accuracy by reducing the standard deviation of the EOBT error to about 2 min, the benefits in terms of taxi-out time savings increases by 15% at EWR, 75% at DFW, and 22% at CLT. This improvement in EOBT accuracy yields an additional fuel burn savings of 600K kg/yr at EWR, 3.4M kg/yr at DFW, and 1.1M kg/yr at CLT. Assuming the jet fuel price to be \$0.68/kg (IATA 2018), the additional annual fuel savings will correspond to \$400K at EWR, \$2.3M at DFW, and \$740K at CLT. The airlines can use these estimates to decide on investment decisions that could improve their EOBT accuracy.

The parametric analysis assumed that the EOBT uncertainty distribution was the same for all airlines at an airport. However, seen in Section 4, the uncertainty in EOBT varies by airline. An interesting question is whether an airline would be rewarded for improved EOBT accuracy by achieving greater taxi-out time reduction. A related question would involve opportunities for airlines to manipulate the system by publishing inaccurate EOBTs to obtain more benefits.

To incentivize airlines to publish more accurate EOBTs, the ATD-2 program categorizes flights into a "planning" group and an "uncertainty" group based on the historical accuracy of their EOBT information. The idea is to give priority to the flights in the planning group so that they are able to get more benefits. Challenges remain with this approach: (a) Deciding on a criterion to categorize flights into two groups (airlines could game by reporting accurate EOBTs during periods with less congestion and incorrect EOBTs during periods with congestion to obtain more benefits); and (b) developing strategies to disincentivize the uncertain group. One possible approach is to penalize airlines if they do not meet the TOBT assigned by the controller. For example, under the Airport Collaborative Decision Making concept, if a flight is unable to pushback within TOBT+5 min, its existing slot is cancelled and it receives a new TOBT (Changi Airport 2018).

5.5. Impact of EOBT information updates

This paper did not consider the impact of EOBT updates by the airlines. As we had discussed in Section 4, airlines update the EOBT information at irregular intervals until a flight pushes back from the gate. If gate-hold decisions are made 20 min prior to a flight's EOBT which continues to be updated, then one would have to determine the most efficient way to accommodate flights with updated EOBTs.

If an airline updates a flight's EOBT, we have one of the following two cases: (a) the new EOBT is prior to the assigned pushback time, in which case, the flight needs to wait at the gate before pushing

back (equivalent to an additional hold); or (b) the new EOBT is after the assigned pushback time. In the latter case, the controller considers the flight to be entering the system at $\max((t + T_p), \text{EOBT})$, and reassigns the gate-hold for that flight. In the above discussion, t is the time at which the EOBT was updated. Both cases penalize the airline, incentivizing airlines to report a more stable EOBT. This is just one approach by which EOBT updates can be accommodated; we could also devise alternative strategies.

5.6. Other considerations

Ramp operations (including clearance for gate pushbacks) in the US are primarily managed by the airlines and not by the FAA. While spot metering has been proposed as a solution to this ‘jurisdictional’ challenge, its practicality remains to be demonstrated. In the departure metering field tests at CLT under the ATD-2 program, gate-hold decisions were made by the ramp controllers (airline personnel) by coordinating with the single airline that handles more than 90% of the operations. The methodology presented in this paper can be adapted to consider spot metering, assuming that the implementation challenges can be addressed.

This paper focused on uncertainties pertaining to EOBTs and taxi times. However, there are multiple sources of uncertainty on the airport surface; for example, departures being subject to Traffic Management Initiatives (TMIs) such as Approval Requests (APREQs) or Miles-In-Trail (MIT) restrictions. However, only a small fraction of flights (fewer than 9% of CLT departures) are affected by such TMIs. We therefore believe that we have considered the uncertainties that have the most significant impact on departure metering.

6. Conclusions

This paper presented an approach to developing queuing network models of airport surface traffic. The proposed models reflected the time-varying nature of capacity and demand, and could be used to predict the queue lengths at congested locations on the airport surface and the taxi times of aircraft. The modeling approach was applied to three U.S airports (EWR, DFW and CLT) with different operating characteristics in order to demonstrate its general applicability.

The queuing network models were used to evaluate two approaches for departure metering to reduce airport taxi-out times: NASA’s ATD-2 logic and a new optimal control algorithm. The proposed queuing network models could be used to estimate the optimal value of the excess queue time buffer, which is an important parameter within the ATD-2 logic. Departure metering algorithms require as inputs information on departure demand, usually in the form of the Earliest Off-Block Times or EOBTs. This paper evaluated the accuracy of the EOBT information currently published by the airlines, and conducted a parametric analysis of the impacts of EOBT uncertainty on the performance of departure metering algorithms. The results showed that EOBT uncertainty can

significantly reduce the departure metering benefits in terms of taxi-out time and fuel burn savings. For example, reducing the standard deviation of the EOBT error from 6 min to 4 min at a 20-min planning horizon would yield an additional annual fuel burn savings of 200 metric tonnes at EWR, 1,300 metric tonnes at DFW, and 800 metric tonnes at CLT. This analysis can potentially incentivize airlines to improve the accuracy of reported EOBTs. Finally, we showed that the proposed optimal control approach to departure metering can handle EOBT uncertainty while still providing taxi-out time reductions.

Acknowledgments

DISTRIBUTION STATEMENT A. Approved for public release. Distribution is unlimited.

This material is based upon work supported by the Federal Aviation Administration under Air Force Contract No. FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Federal Aviation Administration.

References

- Badrinath S, Balakrishnan H, 2017 *Control of a non-stationary tandem queue model of the airport surface. American Control Conference (ACC), 2017*, 655–661 (IEEE), URL <http://dx.doi.org/10.23919/ACC.2017.7963027>.
- Badrinath S, Balakrishnan H, Clemons E, Reynolds TG, 2018 *Evaluating the impact of uncertainty on surface operations. 2018 Aviation Technology, Integration, and Operations Conference (AIAA)*, URL <http://dx.doi.org/10.2514/6.2018-4242>.
- Badrinath S, Li MZ, Balakrishnan H, 2018 *Integrated surface–airspace model of airport departures. Journal of Guidance, Control, and Dynamics* 42(5):1049–1063.
- Ball M, Vossen T, Hoffman R, 2001 *Analysis of demand uncertainty effects in ground delay programs. 4th USA/Europe air traffic management R&D seminar*, 51–60.
- Ball MO, Hoffman RL, Knorr D, Wetherly J, Wambsganss M, 2001 *Assessing the benefits of collaborative decision making in air traffic management. Progress in Astronautics and Aeronautics* 193:239–252.
- Bosson CS, Sun D, 2016 *Optimization of airport surface operations under uncertainty. Journal of Air Transportation* 84–92.
- Burgain P, Pinon OJ, Feron E, Clarke JP, Mavris DN, 2012 *Optimizing pushback decisions to value airport surface surveillance information. IEEE Transactions on Intelligent Transportation Systems* 13(1):180–192, URL <http://dx.doi.org/10.1109/TITS.2011.2166388>.
- Changi Airport, 2018 *Changi Airport A-CDM Handbook*. <http://changiairport-cdm.sg/files/Changi\%20A-CDM\%20Handbook\%20v2.pdf>, retrieved Dec 12, 2018.

- Cheng J, Hoff A, Tittsworth J, Gallo WA, 2016 *The development of wake turbulence re-categorization in the United States. 8th AIAA Atmospheric and Space Environments Conference*, 3434, URL <http://dx.doi.org/10.2514/6.2016-3434>.
- Diana T, 2018 *An evaluation of the impact of wake vortex re-categorization: The case of Charlotte Douglas International airport (CLT). Transportation Research Part A: Policy and Practice* 109:41–49.
- EUROCONTROL Airport CDM Team, 2017 *Airport CDM Implementation–The Manual*. <https://www.eurocontrol.int/sites/default/files/publication/files/airport-cdm-manual-2017.PDF>, retrieved Dec 10, 2018.
- FAA Surface CDM Team, 2012 *US Airport Surface Collaborative Decision Making (CDM) Concept of Operations (ConOps) in the Near-Term: Application of Surface CDM at United States Airports*.
- Federal Aviation Administration, 2018a *Airport Surface Detection Equipment, Model X (ASDE-X)*. https://www.faa.gov/air_traffic/technology/asde-x/, retrieved Dec 14, 2018.
- Federal Aviation Administration, 2018b *Aviation System Performance Metrics (ASPM)*. <http://aspm.faa.gov/>, retrieved Oct 18, 2018.
- Federal Aviation Administration, 2018c *TFDM Overview*. <https://www.faa.gov/>, retrieved Dec 12, 2018.
- Gopalakrishnan K, Balakrishnan H, Jordan R, 2016 *Deconstructing delay dynamics: An air traffic network example. International Conference on Research in Air Transportation (ICRAT)*.
- Hansen M, Nikoleris T, Lovell D, Vlachou K, Odoni A, 2009 *Use of queuing models to estimate delay savings from 4D trajectory precision. Eighth USA/Europe Air Traffic Management Research and Development Seminar*.
- IATA, 2018 *Jet Fuel Price Monitor*. <https://www.iata.org/publications/economics/fuel-monitor/Pages/index.aspx>, retrieved Dec 14, 2018.
- ICAO, 2018 *ICAO Engine Emissions Databank*. <https://www.easa.europa.eu/easa-and-you/environment/icao-aircraft-engine-emissions-databank>, retrieved Dec 14, 2018.
- Jacquillat A, 2012 *A queuing model of airport congestion and policy implications at JFK and EWR*. Ph.D. thesis, Massachusetts Institute of Technology.
- Khadilkar H, Balakrishnan H, 2014 *Network congestion control of airport surface operations. Journal of Guidance, Control, and Dynamics* 37(3):933–940, URL <http://dx.doi.org/10.2514/1.57850>.
- Khadilkar HD, 2013 *Networked control of aircraft operations at airports and in terminal areas*. Ph.D. thesis, Massachusetts Institute of Technology.
- Kivestu PA, 1976 *Alternative methods of investigating the time dependent M/G/k queue*. Ph.D. thesis, Massachusetts Institute of Technology.
- Lee H, 2014 *Airport surface traffic optimization and simulation in the presence of uncertainties*. Ph.D. thesis, Massachusetts Institute of Technology.

- Lee H, Malik W, Zhang B, Nagarajan B, Jung YC, 2015 *Taxi time prediction at Charlotte Airport using fast-time simulation and machine learning techniques. 15th AIAA Aviation Technology, Integration, and Operation (ATIO) Conference, Dallas, TX*, URL <http://dx.doi.org/10.2514/6.2015-2272>.
- Liu Y, Hansen M, Gupta G, Malik W, Jung Y, 2014 *Predictability impacts of airport surface automation. Transportation Research Part C: Emerging Technologies* 44:128–145.
- Lovell DJ, Vlachou K, Rabbani T, Bayen A, 2013 *A diffusion approximation to a single airport queue. Transportation Research Part C: Emerging Technologies* 33:227–237.
- McFarlane P, Balakrishnan H, 2016 *Optimal control of airport pushbacks in the presence of uncertainties. American Control Conference (ACC), 2016*, 233–239 (IEEE).
- Morrison SA, Winston C, 2007 *Another look at airport congestion pricing. American Economic Review* 97(5):1970–1977, URL <http://dx.doi.org/10.1257/aer.97.5.1970>.
- Murça MCR, 2017 *A robust optimization approach for airport departure metering under uncertain taxi-out time predictions. Aerospace Science and Technology* 68:269–277.
- Nakahara A, Reynolds T, White T, Maccarone C, Dunsy R, 2011 *Analysis of a surface congestion management technique at New York JFK airport. 11th AIAA Aviation Technology, Integration, and Operations (ATIO) Conference, including the AIAA Balloon Systems Conference and 19th AIAA Lighter-Than*, 6987.
- Nakahara A, Reynolds TG, 2013 *Estimating current & future system-wide benefits of airport surface congestion management. 10th USA/Europe Air Traffic Management Research and Development Seminar, Chicago*.
- OAG, 2016 *Airport operations data*. <https://www.oag.com>, retrieved Dec 14, 2018.
- Okuniek N, Sparenberg L, 2017 *Opportunities and challenges when implementing trajectory-based taxi operations at European and US CDM airports. Digital Avionics Systems Conference (DASC), 2017 IEEE/AIAA 36th*, 1–10 (IEEE).
- Pujet N, Delcaire B, Feron E, 2000 *Input-output modeling and control of the departure process of busy airports. Air Traffic Control Quarterly* 8(1):1–32, URL <http://dx.doi.org/10.2514/atcq.8.1.1>.
- Pyrgiotis N, Malone KM, Odoni A, 2013 *Modelling delay propagation within an airport network. Transportation Research Part C: Emerging Technologies* 27:60–75.
- Rappaport D, Yu P, Griffin K, Daviau C, 2009 *Quantitative analysis of uncertainty in airport surface operations. 9th AIAA Aviation Technology, Integration, and Operations Conference (ATIO) and Aircraft Noise and Emissions Reduction Symposium (ANERS)*, 6987.
- Ryerson MS, Woodburn A, 2014 *Build airport capacity or manage flight demand? how regional planners can lead american aviation into a new frontier of demand management. Journal of the American Planning Association* 80(2):138–152, URL <http://dx.doi.org/10.1080/01944363.2014.961949>.

- Sharma S, Capps A, Engelland S, Jung Y, 2018 *Operational impact of the baseline integrated arrival, departure, and surface system field demonstration. IEEE/AIAA 37th Digital Avionics Systems Conference (DASC)*, 1–10 (IEEE).
- Simaiakis I, Balakrishnan H, 2014 *Probabilistic modeling of runway inter-departure times. Journal of Guidance, Control, and Dynamics* 37(6):2044–2048.
- Simaiakis I, Balakrishnan H, 2015 *A queuing model of the airport departure process. Transportation Science* 50(1):94–109, URL <http://dx.doi.org/10.1287/trsc.2015.0603>.
- Simaiakis I, Khadilkar H, Balakrishnan H, Reynolds TG, Hansman RJ, 2014 *Demonstration of reduced airport congestion through pushback rate control. Transportation Research Part A: Policy and Practice* 66:251–267, URL <http://dx.doi.org/10.1016/j.tra.2014.05.014>.
- Simaiakis I, Pyrgiotis N, 2010 *An analytical queuing model of airport departure processes for taxi out time prediction. AIAA Aviation Technology, Integration and Operations (ATIO) Conference*, URL <http://dx.doi.org/10.2514/6.2010-9148>.
- Simaiakis I, Sandberg M, Balakrishnan H, 2014 *Dynamic control of airport departures: Algorithm development and field evaluation. IEEE Transactions on Intelligent Transportation Systems* 15(1):285–295, URL <http://dx.doi.org/10.1109/TITS.2013.2278484>.
- Tipper D, Sundareshan MK, 1990 *Numerical methods for modeling computer networks under nonstationary conditions. IEEE Journal on Selected Areas in Communications* 8(9):1682–1695, URL <http://dx.doi.org/10.1109/49.62855>.
- Verma S, Coupe WJ, Lee H, Robeson I, Jung Y, Sharma S, Dulchinos VL, Stevens L, 2018 *Tactical surface metering procedures and information needs for charlotte douglas international airport. International Conference on Applied Human Factors and Ergonomics*, 157–169 (Springer).
- Wang WP, Tipper D, Banerjee S, 1996 *A simple approximation for modeling nonstationary queues. INFOCOM'96. Fifteenth Annual Joint Conference of the IEEE Computer Societies*, volume 1, 255–262 (IEEE), URL <http://dx.doi.org/10.1109/INFCOM.1996.497901>.
- Windhorst RD, Montoya JV, Zhu Z, Gridnev S, Griffin K, Saraf A, Stroiney S, 2013 *Validation of simulations of airport surface traffic with the surface operations simulator and scheduler. Los Angeles, CA, 13th AIAA Aviation Technology, Integration, and Operations Conference*, URL <http://dx.doi.org/10.2514/6.2013-4207>.