

MIT Open Access Articles

Control and optimization algorithms for air transportation systems

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Balakrishnan, Hamsa. "Control and Optimization Algorithms for Air Transportation Systems." Annual Reviews in Control 41 (2016): 39–46 © 2016 International Federation of Automatic Control

As Published: <http://dx.doi.org/10.1016/J.ARCONTROL.2016.04.019>

Publisher: Elsevier BV

Persistent URL: <http://hdl.handle.net/1721.1/114832>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-NonCommercial-NoDerivs License



Control and Optimization Algorithms for Air Transportation Systems

Hamsa Balakrishnan

*Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA 02139, USA.
hamsa@mit.edu.*

Abstract

Modern air transportation systems are complex cyber-physical networks that are critical to global travel and commerce. As the demand for air transport has grown, so have congestion, flight delays, and the resultant environmental impacts. With further growth in demand expected, we need new control techniques, and perhaps even redesign of some parts of the system, in order to prevent cascading delays and excessive pollution.

In this survey, we consider examples of how we can develop control and optimization algorithms for air transportation systems that are grounded in real-world data, implement them, and test them in both simulations and in field trials. These algorithms help us address several challenges, including resource allocation with multiple stakeholders, robustness in the presence of operational uncertainties, and developing decision-support tools that account for human operators and their behavior.

Keywords: Air transportation, congestion control, large-scale optimization, data-driven modeling, human decision processes

1. Introduction

The air transportation system operated nearly 85 million flights worldwide in 2014, serving 6.7 billion passengers and 102 million metric tons of cargo. The Asia-Pacific region served more than a third of these passengers, while Europe and North America served about a quarter each. Emerging markets in the Middle East are experiencing an annual growth in traffic of more than 10% annually [1]. Although there are nearly 42,000 airports worldwide (nearly 20,000 airports in the United States), traffic demand tends to be concentrated at a small number of them: The top 30 airports serve more than one-third of all passengers, while the busiest airports (Chicago O'Hare, Atlanta and Los Angeles) each see more than 700,000 aircraft operations annually [1, 2].

The increasing demand for air traffic operations has further strained this already capacity-limited system, leading to significant congestion, flight delays, and pollution. Domestic flight delays in the

US have been estimated to cost airlines over \$19 billion and the national economy over \$41 billion annually, waste 740 million gallons of jet fuel, and release an additional 7.1 billion kilograms of CO₂ into the earth's atmosphere [3]. The demand for airspace resources is expected to significantly increase in the upcoming decades, and to also include operations of autonomous aircraft [4, 5]. The networked nature of the air transportation system also leads to the propagation of delays from one part of the system to another. To prevent cascading delays and even congestive collapse, there is a need for new analysis techniques and operational strategies for air transportation systems.

The design of algorithms for air transportation, as in the case of most real-world infrastructures, yields a range of multi-objective optimization problems: For example, one would like to improve the efficiency (in terms of reducing total flight delays, fuel burn, delays per passenger, etc.), robustness (that is, minimize the propagation of delays through the system), while still maintaining the safety and se-

curity of the system. These objectives are difficult to achieve in practice, due to the challenges posed by the presence of uncertainties, human factors, and competing stakeholder interests. However, it is possible to overcome these challenges by leveraging the increasingly available operational data to build simple yet realistic models, and to use these models to develop and implement scalable control and optimization algorithms to improve system performance.

In this paper, we present three examples of how the challenges mentioned above can be addressed in the context of air transportation systems:

1. Airport congestion control.
2. Large-scale optimization algorithms for air traffic flow management.
3. Learning models of air traffic controller decision processes and the associated utility functions.

This paper is based on a semi-plenary lecture given by the author at the American Control Conference, Chicago, IL, 2015.

2. Airport congestion control

Taxiing aircraft consume fuel, and emit pollutants such as Carbon Dioxide, Hydrocarbons, Nitrogen Oxides, Sulfur Oxides and Particulate Matter that impact the local air quality at airports [6, 7, 8, 9]. Although fuel burn and emissions are approximately proportional to the taxi times of aircraft, other factors such as the throttle settings, number of engines that are powered, and pilot and airline decisions regarding engine shutdowns during delays also influence them [10]. Domestic flights in the United States emit about 6 million metric tonnes of CO₂, 45,000 tonnes of CO, 8,000 tonnes of NO_x, and 4,000 tonnes of HC taxiing out for takeoff; almost half of these emissions are at the 20 most congested airports in the country [11]. Aircraft in Europe have been estimated to spend 10-30% of their flight time taxiing [12]. Data also show that 20% of delays at major US airports occur not due to bad weather, but due to high traffic volume [13]. Better congestion management at airports has the potential to mitigate these impacts.

2.1. Impacts of airport congestion

Pujet et al. analyzed surface congestion by considering the takeoff rate of an airport as a function

of the number of aircraft taxiing out [14]. Fig. 1 shows a similar analysis for Philadelphia International Airport (PHL) in 2007, for one runway configuration (set of active runways at the time), under visual meteorological conditions (VMC) [10].

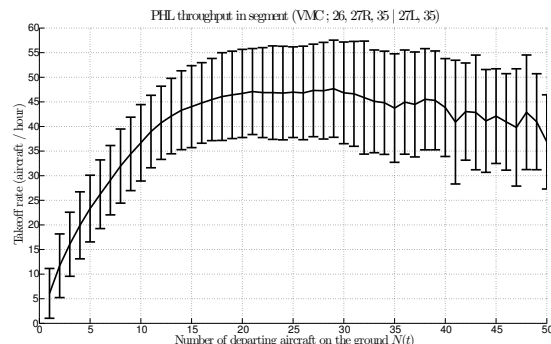


Figure 1: Average take-off rate as a function of the number of departing aircraft on the ground at PHL. The error bars represent the standard deviation of the take-off rate [10].

Fig. 1 illustrates that although the take-off rate increases at first, it saturates once there are approximately 20 departing aircraft on the ground. Any further pushbacks will just lead to congestion, and will not result in an improvement in the take-off rate. It is also worth noting that for a very high numbers of departures on the ground (more than 30 in Fig. 1), the departure throughput can even *decrease* due to surface gridlock. Similar phenomena have been observed at several major airports in the US, including Boston Logan International Airport (BOS), Newark Liberty International Airport (EWR), New York John F. Kennedy International Airport (JFK), New York La Guardia International Airport (LGA), and Charlotte Douglas International Airport (CLT) [10, 15, 16, 17]. This phenomenon of throughput saturation is also typical of queuing systems, motivating the development of queuing network models of major airports [17, 18].

2.2. Congestion management strategies

One of the earliest efforts at airport congestion control was the Departure Planner project [19]. This project proposed the concept of a virtual departure queue, where aircraft would be held (at their gates) until an appropriately determined pushback time. The resultant *N-Control* strategy was a threshold heuristic, where if the total number of departing aircraft on the ground exceeded a certain threshold, N_{ctrl} , any further aircraft requesting pushback

were held at their gates until the number of departures on the ground fell below the threshold [19, 11]. Other variants and extensions of this policy have also been studied [20, 21, 22, 14]. Interestingly, a similar heuristic has been known to be deployed by Air Traffic Controllers at BOS during times of extreme congestion [23]. The N-Control policy is similar in spirit to *constant work-in-process* or CON-WIP policies that have been proposed for manufacturing systems [24].

Several other approaches to departure metering have been proposed, including the Ground Metering Program at New York’s JFK airport [25, 26], the field-tests of the Collaborative Departure Queue Management concept at Memphis (MEM) airport [27], the human-in-the-loop simulations of the Spot and Runway Departure Advisor (SARDA) concept at Dallas Fort Worth (DFW) airport [28], and the trials of the Departure Manager (DMAN) concept [29] at Athens International airport (ATH) [30]. In addition, Mixed Integer Linear Programming (MILP) formulations of surface traffic optimization have been considered, but are generally known to be NP-hard [31, 32, 33, 34, 35]. In practice, these strategies are treated as open-loop policies that are periodically reoptimized. Full-state feedback policies have also been proposed, but have presented practical challenges [36].

2.3. Design and implementation of a congestion control algorithm

While there has been prior research on the optimal control of queuing systems [37, 38], the application of these techniques to airport operations has remained a challenge. In particular, the need to interface with current air traffic control procedures, and the different sources of uncertainty (the variability in departure throughput and the randomness of taxi-out times) pose practical concerns.

2.3.1. Rate control strategies

On-off or event-driven pushback control policies (such as a threshold heuristic) are not desirable in practice, since both air traffic controllers and airlines prefer a pushback rate that is periodically updated. This observation motivates the development of *Pushback Rate Control* policies, wherein an optimal pushback rate is recommended to air traffic controllers for each 15-minute interval, and the rate is updated periodically [11, 39]. The threshold N-Control policy can be adapted to obtain a rate control policy by predicting the average throughput

under saturation over the next 15-minute interval, and then determining the number of pushbacks in that interval that would help maintain the desired level of traffic (typically around the threshold value at which the throughput saturates). Such an approach has been developed and tested at BOS in 2010 with promising results: During the course of eight 3-hour periods, a total of nearly 16 hours of taxi-out time savings were achieved, resulting in fuel burn savings of 10,500-13,000 kg [11]. While simple and easy to implement, this approach does not explicitly consider the variability in throughput and taxi-out times, thereby increasing the risk of runway starvation.

2.3.2. Dynamic programming for Pushback Rate Control

A better approach to accommodating the uncertainties in throughput and taxi-out times is through the formulation of a dynamic control problem. Using a queuing model of the departure process built from operational data, we can use dynamic programming to determine the optimal pushback rate that minimizes taxi-out times, while still maintaining runway utilization [39].




Figure 2: An illustration of the optimal pushback rate as a function of the number of aircraft in the departure queue and the number of aircraft traveling to the runway [39].

The dynamic programming formulation considers the system state consisting of the number of aircraft taxiing to the runway and the length of the departure runway queue. Using predictions of the runway throughput in the next time period, the runway queue is modeled as a semi-Markov process, and the system state is projected by solving the resultant Chapman-Kolmogorov equations. The optimal policy is determined by solving the Bellman equation for the infinite horizon average cost problem [39]. Fig. 2 illustrates one such optimal control policy that is determined for the case of a par-

particular runway configuration at BOS under visual conditions. In comparing the resultant policy with the adapted N-Control policy described in Section 2.3.1, the dynamic programming approach is found to handle uncertainty better (as expected), and results in more robust policies that reduce the risk of runway starvation. The implementation of the dynamic programming based pushback rate control policies at BOS in 2011 showed that during eight 4-hour tests, taxi-out time was reduced by nearly 13 hours, while fuel use was reduced by more than 8,200 kg [39].

2.3.3. Interfacing with air traffic controllers

The optimal pushback rates need to be communicated to the air traffic controllers in the airport tower, with minimal distraction from their responsibilities. For this reason, vocal communications with tower personnel are not desired. We therefore developed Android™ tablet-based decision-support displays to present a color-coded suggested pushback rate (as was done using cards at BOS in 2010), and an alternate display that provided additional support to the controllers. This decision support tool was used to implement the dynamic programming policy at BOS airport in 2011, with positive feedback from the air traffic controllers [40]. Fig. 3 (left) shows a color-coded cards in the tower, while Fig. 3 (right) shows the tablet-based display.

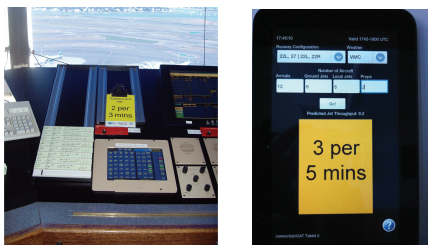


Figure 3: (Left) Display of a pushback rate control card in the Boston airport tower; (right) Tablet-based rate control input interface [40].

2.3.4. Results, extensions and open problems

Over the course of 15 metering periods during the 2010-2011 trials, Pushback Rate Control strategies were found to result in a total fuel savings of 20,800-23,600 kg. The average metered flight was held at the gate for only an additional 4.7 min, and saved more than 52 kg of fuel as a result. The policies were shown to be fair, in that for every minute

of gate-hold that an airline experiences, it also receives a minute of taxi-out time savings. In addition, the policies were shown to accommodate practical constraints, such as gate-use conflicts, when a departure would need to leave the virtual departure queue (its gate) early because the next aircraft to use that gate had arrived [39].

An alternate approach to airport surface congestion management is through drawing an analogy to more general network congestion management problems. The airport is modeled as a network consisting of major taxiways and their intersections, using surface surveillance data [41]. Terminal-area operations and aircraft arrivals can also be accommodated by these models. Although the resultant models are of significantly higher complexity than the ones considered in Section 2.3.2, the optimal control problems can be solved efficiently using approximate dynamic programming [42]. This approach effectively accounts for operational uncertainties, and practical resource constraints such as limited gate availability. Integrated arrival-departure control policies can be shown to yield taxi-out time and fuel burn reductions of 10% across all flights operating at an airport, while reducing practical problems such as gate conflicts by 30% [42].

Finally, the wide-spread deployment of departure metering strategies requires the adaptation of these algorithms to a range of airport operating environments [16, 43]. In order to do so, several open research questions, such as the impact of uncertainty on the efficacy of departure metering, as well as the value of information-sharing among different stakeholders, need to be studied. These are topics of ongoing research on airport surface operations.

3. Large-scale, distributed air traffic flow management

Air Traffic Flow Management (ATFM) is the process of modifying departure times and trajectories of flights in order to address congestion, namely imbalances between available resource capacity and demand, and thereby reduce delay costs. These adjustments are typically made strategically, a few hours ahead of flight operations. Capacity-demand imbalances can occur either because capacity is reduced (for example, due to weather impacts) or because demand is high (for example, over-scheduling during peak periods). Weather is estimated to cause nearly two-thirds of flight delays, while high traffic

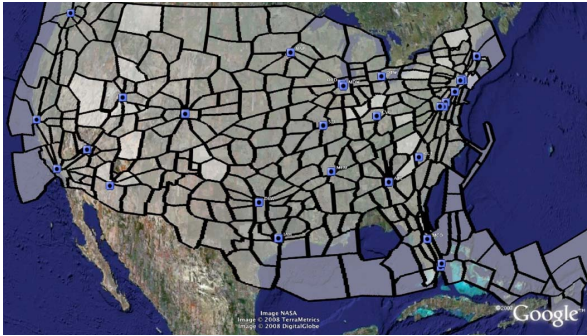


Figure 4: Sector boundaries for enroute sectors in the continental US. The markers denote the top 30 airports [44].

demand is responsible for nearly 20% of delays at major US airports [13].

The air transportation system consists of many interconnected resources, the chief among which are the airports and airspace sectors. Fig. 4 shows the 30 major US airports, along with the high altitude enroute sectors [44]. Although airports are typically the most constrained air traffic resources in the US, airspace sectors may also experience congestion. By contrast, airspace sector congestion is a more frequent problem in Europe [45].

There are two main challenges to the development of ATFM algorithms: First, weather influences capacity, and tends to be uncertain and dynamic in nature, requiring algorithms that can be easily updated in the event of new information; and secondly, flight connectivity implies that the same aircraft may operate multiple flights in a day resulting in delay propagation. Nearly one-third of all domestic flight delays in the US are because the previous flight operated by that aircraft arrived late [46]. Fig. 5 shows the flight connectivity on a typical day in the US: Only about 6% of flights have no connection, while aircraft typically operate 4–6 flights in a day. Such high levels of connectivity make myopic, rolling horizon formulations of ATFM significantly suboptimal.

The control actions available in determining a flight trajectory are ground delays (i.e., to delay the departure of the aircraft so as to arrive at a constrained resource at a different time), airborne delays (i.e., modify its speed), rerouting (i.e., changing its spatial path), and cancellation (i.e., not operate the flight on that day). Of these options, ground delays have the lowest cost per unit time (since they are absorbed on the ground, and frequently

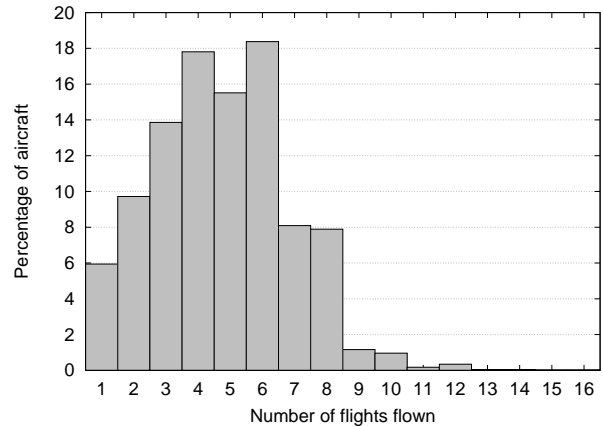


Figure 5: Histogram showing the level of connectivity on July 8, 2013 with a total of 4,054 distinct aircraft and 19,217 flights (average of 4.75 flights per aircraft) [47]. Data from [46].

while the aircraft is parked at the gate with engines off), airborne delays are more expensive than ground delays (since the aircraft is active and in the air). Ground delays are also considered “safer” than airborne delays. A reroute requires additional coordination among stakeholders, and can be expensive if it is significantly longer than the nominal route. A cancellation is the most expensive option, since it implies that the passengers and crew must be reaccommodated on other flights. In addition, connectivity implies that if a flight is significantly delayed or cancelled early in the day, all subsequent operations by that aircraft may have to be delayed or cancelled. The goal of ATFM is to maximize a system objective which is the sum of flight-specific objectives; where each flight achieves some benefit (revenue) when it is operated, but also incurs a cost depending on the trajectory flown.

Nearly every element in the air transportation system is capacity-limited. As a result, ATFM algorithms are faced with the task of not just determining the trajectory of each aircraft (i.e., the departure time, route, enroute speeds along different segments of its route, etc. for each flight that it operates over the course of a day), but also need to satisfy the capacity limits and other operational constraints described below:

1. *Airspace sector capacity constraints:* These constraints limit the number of aircraft that can be in a given airspace sector at any time. The actual values depend on the size and geometry of the sector, as well as air traffic con-

troller workload [48, 49].

2. *Airport capacity constraints:* These constraints limit the arrival and departure throughputs at an airport at any time. Since airport resources such as runways are shared by arrivals and departures, airport capacities are represented as envelopes that represent the trade-off between arrival and departure capacities at any time [50, 51]. An example of such a capacity envelope for Newark (EWR) airport is shown in Fig. 6. The capacity envelope is usually modeled as a polytope, and depends on the airport, choice of runway configuration, weather conditions, etc.




Figure 6: Observed capacity envelope for Newark airport, under good weather conditions [51]. Any pair of arrival/departure rates within the envelope is achievable.

3. *Operational constraints:* Operational constraints are limitations on what actions may be performed by a specific flight. They include minimum and maximum transit times on airspace links, maximum ground and airborne delay that can be incurred by a flight, minimum turnaround time between successive flights for an aircraft, and any routing restrictions. These constraints may also vary by aircraft performance characteristics such as the nominal speed and altitude.

The deterministic ATFM problem can be described as follows:

Given a set of flights (and associated aircraft operating them), and airport and airspace capacity constraints, identify a trajectory for each aircraft that maximizes a system-wide objective (difference between benefit and cost, summed over all flights), and that obeys operational and capacity constraints for all time periods.

3.1. Algorithms for ATFM

The problem of developing automation and decision support for ATFM has been a rich topic of research for several decades [52, 53, 49]. The ATFM problem has typically been formulated as a very large-scale integer program, and has been shown to be NP-hard [54]. Extensions that consider limited routing and rerouting have also been considered [55, 56]. Most of these prior approaches have faced computational challenges (for example, the prior state-of-the-art considered problems with $\sim 6,745$ flights, 30 airports and 145 sectors; a time-discretization of 15-minutes and a planning time horizon of 8 hours; with a computation time of approximately 10 minutes [56]).

In order to address these computational challenges, researchers have also developed models that do not consider space-time trajectories for each aircraft, but instead consider aggregate flows [57, 58, 59, 60, 61]. Eulerian models have been shown to be able to have reasonable predictive capabilities [59], and to reflect the current manner in which air traffic controllers conduct handoffs of aircraft between airspace sectors [82]. Eulerian models have been found to be amenable to the development of feedback control schemes, in a centralized setting [58, 62], in a decentralized setting for networks with a single origin and destination [63], and in distributed multi-airport network settings [64]. These techniques have also been shown to guarantee stability of aircraft queues in each sector, thereby better managing air traffic controller workload [64]. The state-of-the-art in Eulerian models has considered only airborne delays (no ground delays) with 3,419 flight-paths and 284 sectors, at a time-discretization of 1-minute and a planning time horizon of 2 hours, and achieved run times of approximately 21 minutes [65]. While computationally more tractable than the disaggregate models, these Eulerian models are of considerably lower fidelity than the integer programming based ones, and do not model individual trajectories.

3.2. Large-scale, optimal ATFM

In recent work, we have developed a new algorithm to solve very large-scale ATFM problems in a fast and scalable manner. Given flight-specific operating and delay costs, our method determines optimal trajectories, taking into account network and flight connectivity constraints as well as uncertain airport and airspace capacities [47]. Using a column

450 generation based formulation, the problem can be
 efficiently decomposed into a set of parallelizable
 sub-problems, with an easy-to-solve master problem 485
 that coordinates between the sub-problems. Fig.
 7 shows a schematic of the solution process [47].

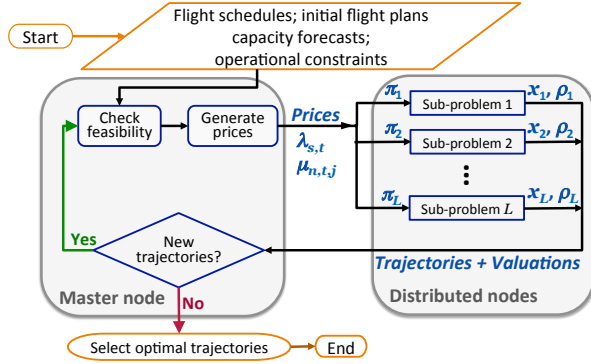


Figure 7: Schematic of the distributed solution process for very large-scale ATFM [47].

455 3.2.1. Results, extensions and open problems

Computational experiments using US nation-scale
 examples drawn from operational data show that
 the proposed approach can determine very good
 (close-to-optimal) integer solutions. For instances
 with $\sim 17,500$ flights, 370 airports and 375 sectors,
 at a time-discretization of 5-minutes and a planning
 time horizon of 24 hours, the computation time is
 found to be under 5 min, a significant improvement
 over prior state-of-the-art [47].

465 The easily parallelizable nature of the approach,
 in addition to having computational benefits, has
 the potential to enable distributed, yet collabora-
 tive, decision-making among the different airlines
 [66]. In order for any resource allocation process
 such as ATFM to yield efficient outcomes, airlines
 must be incentivized to participate, and to also
 truthfully report their delays and cancellations.
 These issues have received only limited attention
 to date, and only in the context of single-resource
 allocation [67, 68] or for aggregate models [69, 70].
 The analysis of these issues in the context of networked
 ATFM problems remains an important open challenge.

480 4. Determining utility functions of human decision processes

Like most modern infrastructure systems, the
 performance of the air transportation system de-

490 pends significantly on decisions made by human
 operators. Modeling these systems and providing
 decision support to operators needs an understand-
 ing of the objective functions in the decision pro-
 cesses, in addition to efficient algorithms that can
 optimize them. Currently, the use of idealized ob-
 jective functions (that do not reflect the true sys-
 tem goals) and the difficulty in adapting decision
 support tools to particular operating environments
 (which can take months, or even years) pose signifi-
 cant barriers to implementing advanced algorithms.
 For these reasons, the problem of *inverse optimiza-
 tion*, or reverse-engineering the objective functions
 that best reflect the decision-maker's desire, is an
 important one. We consider this problem for the
 case of airport configuration selection.

Most major airports possess multiple runways
 (Fig. 8), and a subset of these runways (and as-
 sociated traffic directions) are selected at any time
 to handle arrivals and departures. This choice of
 runways is known as the airport or runway confi-
 guration, and is a strong driver of the airport capac-
 ity envelope. As seen in Section 3, airport capacity
 is an essential input to ATFM algorithms. Several
 factors, including weather conditions (wind and vis-
 ibility), traffic demand, air traffic controller work-
 load, and the coordination of flows with neighboring
 airports influence the selection of runway confi-
 guration. However, little is known about the relative
 weightings given to the different factors that influ-
 ence runway configuration selection.

4.1. Runway configuration selection

495 Two classes of models have been developed for
 runway configuration selection: Prescriptive mod-
 els and descriptive models. Most prior research has
 belonged to the former class, and aim to recom-
 mend an optimal runway configuration, subject to
 operational constraints. These include efforts to op-
 timally schedule runway configurations, taking into
 account different models of weather forecasts and
 the loss of capacity during configuration switches
 [71, 72, 73, 74, 75].

495 Descriptive models analyze historical data in or-
 der to predict the runway configuration selected by
 the decision-makers, and have received limited at-
 tention. A 24-hour forecast of runway configura-
 tion was developed for Amsterdam Schiphol airport,
 using a probabilistic weather forecast [76]. A logistic
 regression-based approach was used to develop a
 descriptive model of runway configuration selection

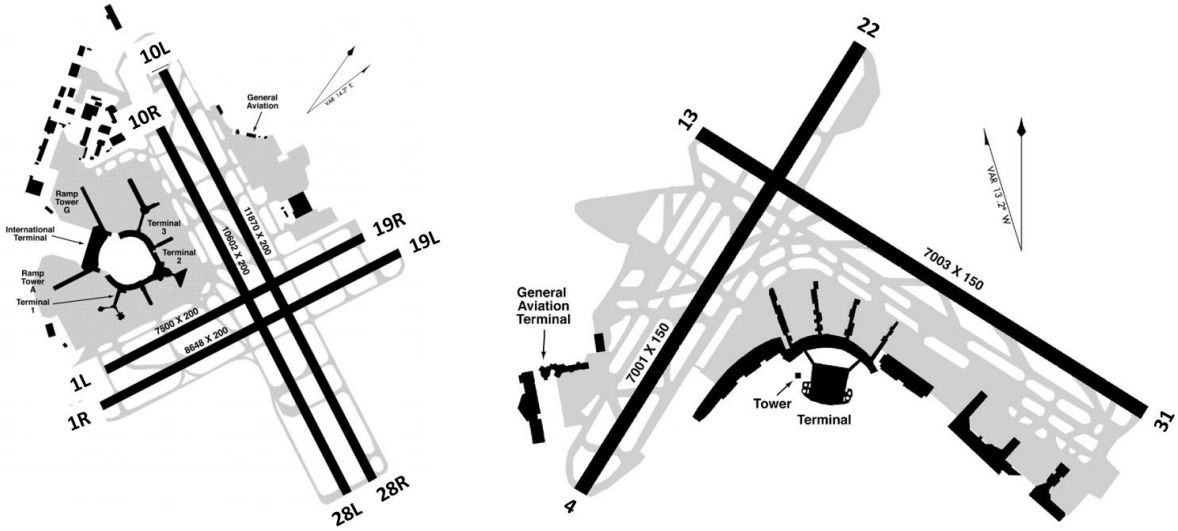


Figure 8: Layouts of San Francisco (SFO) and La Guardia (LGA) Airports.

at LaGuardia (LGA) and John F. Kennedy (JFK) airports [77].

4.2. Discrete-choice models

Discrete-choice models are behavioral models that describe the choice selection of a decision maker, or the nominal decision selection among an exhaustive set of possible alternative options, called the choice set [78]. Each alternative in the choice set is assigned a utility function based on defining attributes that are related to the decision selection process. At any given time, the feasible alternative with the maximum utility is assumed to be selected by the decision maker.

In other words, the utility function is modeled as stochastic random variable, with an observed (deterministic) component, V , and a stochastic error component, ϵ . For the n^{th} selection, given a set of feasible alternatives C_n , the utility of choice $c_i \in C_n$ is represented as

$$U_{n,i} = V_{n,i} + \epsilon_{n,i}. \quad (1)$$

The decision maker selects the alternative j with maximum utility, that is, $c_j \in C_n$ that maximizes $U_{n,j}$. The random error component of the utility function reflects all measurement errors, including unobserved attributes, variations between different decision-makers, proxy variable effects, and reporting errors. The Gumbel distribution is used to approximate a normal distribution due to its

computational advantages. Different model structures (Multinomial Logit, Nested Logit, etc.) correspond to different assumptions on the correlations between the error terms [78, 79, 80].

4.3. Results, extensions and open problems

Maximum-likelihood estimates of the linear utility functions and the underlying structure can be estimated from the training data. The estimation problem is a nonlinear optimization problem, and is solved computationally using an open-source software package called BIOGEME [81]. Case studies from Newark (EWR), LaGuardia (LGA) and San Francisco (SFO) airports have demonstrated an over 20% improvement in the predictions of actual runway configuration selection decisions compared to prior models [79, 80].

Conclusions

The objective of the three research vignettes discussed in this survey paper was to demonstrate the value of real-world operational data in the development of control and optimization algorithms for air transportation systems. In particular, we see that such approaches have the potential to enhance system efficiency, robustness and safety, while addressing the challenges presented by uncertainty, human operators and competition.

Acknowledgments

The work presented in this survey would not be possible without the help of many of my students, collaborators and colleagues, especially Jacob Avery (MIT), Bala Chandran (Resilient Ops, Inc.), Eric Feron (Georgia Tech), John Hansman (MIT), Harshad Khadilkar (TCS Labs), Hanbong Lee (NASA), Varun Ramanujam (Google), Tom Reynolds (MIT Lincoln Lab), Melanie Sandberg (MIT Lincoln Lab), Ioannis Simaiakis (McKinsey & Co.), and Claire Tomlin (UC Berkeley). The research described in Sections 2 and 4 was supported in part by the Federal Aviation Administration and the National Science Foundation. Any opinions stated in this paper are those of the authors alone, and not the funding agencies.

References

- [1] Airports Council International, 2014 World Airport Traffic Report (2015).
- [2] Central Intelligence Agency, CIA World Factbook (2015).
- [3] Joint Economic Committee, US Senate, Your Flight has Been Delayed Again: Flight Delays Cost Passengers, Airlines, and the US Economy Billions (2008).
- [4] Joint Planning and Development Office, Next Generation Air Transportation System Integrated Plan, http://www.jpdo.aero/integrated_plan.html (December 2004).
- [5] United States Government Accountability Office, Unmanned Aerial Systems: FAA Continues Progress toward Integration into the National Airspace, <http://www.gao.gov/assets/680/671469.pdf> (July 2015).
- [6] K. N. Yu, Y. P. Cheung, T. Cheung, R. C. Henry, Identifying the impact of large urban airports on local air quality by nonparametric regression, *Atmospheric Environment* 38 (27) (2004) 4501–4507.
- [7] D. C. Carslaw, S. D. Beevers, K. Ropkins, M. C. Bell, Detecting and quantifying aircraft and other on-airport contributions to ambient nitrogen oxides in the vicinity of a large international airport, *Atmospheric Environment* 40 (28) (2006) 5424–5434.
- [8] M. A. Miracolo, C. J. Hennigan, M. Ranjan, N. T. Nguyen, T. D. Gordon, E. M. Lipsky, A. A. Presto, N. M. Donahue, A. L. Robinson, Secondary aerosol formation from photochemical aging of aircraft exhaust in a smog chamber, *Atmospheric Chemistry and Physics* 11 (9) (2011) 4025–4610.
- [9] A. Ashok, I. C. Dedoussi, S. H. L. Yim, H. Balakrishnan, S. R. H. Barrett, Quantifying the air quality-CO₂ tradeoff potential for airports, *Atmospheric Environment* 99 (2014) 546–555.
- [10] I. Simaiakis, H. Balakrishnan, Analysis and control of airport departure processes to mitigate congestion impacts, *Transportation Research Record: Journal of the Transportation Research Board* (2010) 22–30.
- [11] I. Simaiakis, H. Khadilkar, H. Balakrishnan, T. G. Reynolds, R. J. Hansman, Demonstration of Reduced Airport Congestion through Pushback Rate Control, *Transportation Research Part A: Policy and Practice* 66 (2014) 251–267.
- [12] C. Cros, C. Frings, Alternative taxiing means – Engines stopped, Presented at the Airbus workshop on “Alternative taxiing means – Engines stopped” (November 2008).
- [13] Federal Aviation Administration (FAA), OPSNET website, <http://aspm.faa.gov/> (Accessed 2015).
- [14] N. Pujet, B. Delcaire, E. Feron, Identification and control of the departure process at busy airports, *Air Traffic Control Quarterly* 8 (1).
- [15] I. Simaiakis, Analysis, Modeling and Control of the Airport Departure Process, Ph.D. thesis, Massachusetts Institute of Technology (2012).
- [16] M. Sandberg, T. G. Reynolds, H. Khadilkar, H. Balakrishnan, Demonstration of reduced airport congestion through Pushback Rate Control, in: 10th Eurocontrol/FAA ATM R&D Seminar, 2013.
- [17] I. Simaiakis, H. Balakrishnan, A Queuing Model of the Airport Departure Process, *Transportation Science. Articles in Advance*.
- [18] A. Jacquillat, A queuing model of airport congestion and policy implications at JFK and EWR, Master’s thesis, Massachusetts Institute of Technology (2012).
- [19] E. R. Feron, R. J. Hansman, A. R. Odoni, R. B. Cots, B. Delcaire, W. D. Hall, H. R. Idris, A. Muharremoglu, N. Pujet, The Departure Planner: A conceptual discussion, Tech. rep., Massachusetts Institute of Technology (1997).
- [20] P. Burgain, E. Feron, J.-P. Clarke, A. Darrasse, Collaborative Virtual Queue: Fair Management of Congested Departure Operations and Benefit Analysis, Arxiv preprint arXiv:0807.0661.
- [21] F. Carr, A. Evans, E. Feron, J.-P. Clarke, Software tools to support research on airport departure planning, in: Digital Avionics Systems Conference, IEEE, Irvine CA, 2002.
- [22] F. Carr, Stochastic modeling and control of airport surface traffic, Master’s thesis, Massachusetts Institute of Technology (2001).
- [23] R. Clewlow, D. Michalek, Logan Control Tower: Controller Positions, Processes, and Decision Support Systems, Tech. rep., Massachusetts Institute of Technology (2010).
- [24] M. Spearman, M. Zazanis, Push and pull production systems: Issues and comparisons, *Operations research* (1992) 521–532.
- [25] A. Nakahara, T. Reynolds, T. White, R. Dunskey, Analysis of a Surface Congestion Management Technique at New York JFK Airport, in: AIAA Aviation Technology, Integration and Operations Conference, IEEE, Virginia Beach, VA, 2011.
- [26] S. Stroiney, B. Levy, H. Khadilkar, H. Balakrishnan, Assessing the impacts of the JFK Ground Metering Program, in: IEEE Digital Avionic Systems Conference (DASC), IEEE, 2013.
- [27] C. Brinton, C. Provan, S. Lent, T. Prevost, S. Passmore, Collaborative Departure Queue Management: An Example of Collaborative Decision Making in the United States, in: 9th USA/Europe Air Traffic Management Research and Development Seminar (ATM2011), Berlin, Germany, 2011.

- [28] Y. Jung, T. Hoang, J. Montoya, G. Gupta, W. Malik, L. Tobias, Performance Evaluation of a Surface Traffic Management Tool for Dallas/Fort Worth International Airport, in: 9th USA/Europe Air Traffic Management Research and Development Seminar (ATM2011), Berlin, Germany, 2011.
- [29] D. Böhme, Tactical departure management with the Eurocontrol/DLR DMAN, in: 6th USA/Europe Air Traffic Management Research and Development Seminar, Baltimore, MD, 2005.
- [30] M. Schaper, G. Tsoukala, R. Stavrati, N. Papadopoulos, Departure flow control through takeoff sequence optimisation: Setup and results of trials at Athens airport, in: 2011 IEEE/AIAA 30th Digital Avionics Systems Conference (DASC), IEEE, 2011, pp. 2B2-1.
- [31] J. W. Smeltink, M. J. Soomer, P. R. de Waal, R. D. van der Mei, An Optimisation Model for Airport Taxi Scheduling, in: Thirtieth Conference on the Mathematics of Operations Research, Lunteren, The Netherlands, 2005.
- [32] S. Rathinam, J. Montoya, Y. Jung, An optimization model for reducing aircraft taxi times at the Dallas Fort Worth International Airport, in: 26th International Congress of the Aeronautical Sciences, Anchorage, AK, 2008.
- [33] H. Balakrishnan, Y. Jung, A Framework for Coordinated Surface Operations Planning at Dallas Fort Worth International Airport, in: AIAA Guidance, Navigation, and Control Conference, Hilton Head, SC, 2007.
- [34] H. Lee, H. Balakrishnan, Optimization of Airport Taxiway Operations at Detroit Metropolitan Airport (DTW), in: AIAA Aviation Technology, Integration and Operations Conference, 2010.
- [35] H. Lee, Airport Surface Traffic Optimization and Simulation in the Presence of Uncertainties, Ph.D. thesis, Massachusetts Institute of Technology (2014).
- [36] P. Burgain, O. Pinon, E. Feron, J.-P. Clarke, D. Mavris, On the value of information within a collaborative decision making framework for airport departure operations, in: Digital Avionics Systems Conference, IEEE, 2009.
- [37] T. Crabill, D. Gross, M. Magazine, A classified bibliography of research on optimal design and control of queues, *Operations Research* 25 (2) (1977) 219–232.
- [38] S. Stidham, R. Weber, A survey of Markov decision models for control of networks of queues, *Queueing Systems* 13 (1) (1993) 291–314.
- [39] I. Simaiakis, M. Sandberg, H. Balakrishnan, Dynamic control of airport departures: Algorithm development and field evaluation, *IEEE Transactions on Intelligent Transportation Systems* 15 (1) (2014) 285–295.
- [40] M. Sandberg, I. Simaiakis, H. Balakrishnan, T. Reynolds, R. Hansman, A decision support tool for the pushback rate control of airport departures, *IEEE Transactions on Human-Machine Systems* 44 (3) (2014) 416–421.
- [41] H. Khadilkar, H. Balakrishnan, Network congestion control of airport surface operations, *Journal of Guidance, Control and Dynamics* 37 (3) (2014) 933–940.
- [42] H. Khadilkar, H. Balakrishnan, Network congestion control of airport surface operations, *Integrated Control of Airport and Terminal Airspace Operations* 24 (1) (2016) 216–225.
- [43] Federal Aviation Administration, NextGen Implementation Plan (March 2011).
- [44] J. L. Ny, H. Balakrishnan, Distributed Feedback Control for an Eulerian Model of the National Airspace System, in: American Control Conference (ACC), St. Louis, MI, 2009.
- [45] G. Lulli, A. Odoni, The European Air Traffic Flow Management Problem, *Transportation Science* 41 (4) (2007) 431–443.
- [46] Bureau of Transportation Statistics, Airline On-Time Statistics and Delay Causes, <http://www.transtats.bts.gov/> (2015).
- [47] H. Balakrishnan, B. G. Chandran, Optimal large-scale air traffic flow management, www.mit.edu/hamsa/pubs/BalakrishnanChandran_ATFM.pdf (2014).
- [48] Federal Aviation Administration, Facility Operation and Administration: Monitor Alert Parameter (JO 7210.3W), <http://tfmlearning.faa.gov/Publications/atpubs/FAC/1708.html> (2011).
- [49] T. W. M. Vossen, R. Hoffman, A. Mukherjee, Air Traffic Flow Management, in: C. Barnhart, B. C. Smith (Eds.), *Quantitative Problem Solving Methods in the Airline Industry*, Vol. 169 of International Series in Operations Research & Management Science, Springer Science+Business Media, LLC, 2012, pp. 385–453.
- [50] E. P. Gilbo, Airport capacity: Representation, estimation, optimization, *IEEE Transactions on Control Systems Technology* 1 (3) (1993) 144–154.
- [51] U.S. Department of Transportation, Airport capacity benchmark 2004, Tech. rep. (2004).
- [52] A. R. Odoni, The flow management problem in air traffic control, in: A. R. Odoni, L. Bianco, G. Szego (Eds.), *Flow Control of Congested Networks*, Springer-Verlag, Berlin, 1987, pp. 269–288.
- [53] D. Bertsimas, A. Odoni, A critical survey of optimization models for tactical and strategic aspects of air traffic flow management, Tech. Rep. CR-97-206409, NASA (1997).
- [54] D. Bertsimas, S. Stock Patterson, The air traffic flow management problem with enroute capacities, *Operations Research* 46 (3) (1998) 406–422.
- [55] D. Bertsimas, S. Stock Patterson, The traffic flow management rerouting problem in air traffic control: A dynamic network flow approach, *Transportation Science* 34 (3) (2000) 239–255.
- [56] D. Bertsimas, G. Lulli, A. Odoni, An integer optimization approach to large-scale air traffic flow management, *Operations Research* 59 (1) (2011) 211–227.
- [57] P. K. Menon, G. D. Sweriduk, K. D. Bilimoria, New approach for modeling, analysis, and control of air traffic flow, *Journal of Guidance, Control, and Dynamics* 27 (5) (2004) 737–744.
- [58] P. K. Menon, G. D. Sweriduk, T. Lam, G. M. Diaz, K. Bilimoria, Computer-aided Eulerian air traffic flow modeling and predictive control, *AIAA Journal of Guidance, Control and Dynamics* 29 (2006) 12–19.
- [59] D. Sun, I. S. Strub, A. Bayen, Comparison of the performance of four Eulerian network flow models for strategic air traffic management, *Networks and Heterogeneous Media* 2 (4) (2007) 569–594.
- [60] S. Roy, B. Sridhar, G. C. Verghese, An aggregate dynamic stochastic model for air traffic control, in: *Proceedings of the 5th USA/Europe ATM 2003 R&D Seminar*, Budapest, Hungary, 2003.

- [61] B. Sridhar, T. Soni, K. Sheth, G. Chatterji, An aggregate flow model for air traffic management, in: AIAA Conference on Guidance, Navigation, and Control, no. AIAA Paper 2004-5316, Providence, RI, 2004.
- [62] A. Mukherjee, M. Hansen, A dynamic rerouting model for air traffic flow management, *Transportation Research Part B: Methodological* 43 (1) (2009) 159 – 171.
- [63] H. Arneson, C. Langbort, Distributed sliding mode control design for a class of positive compartmental systems, in: *Proceedings of the American Control Conference*, 2008.
- [64] J. L. Ny, H. Balakrishnan, Feedback control of the national airspace system, *Journal of Guidance, Control, and Dynamics* 34 (3) (2011) 832–846.
- [65] P. Wei, Y. Cao, D. Sun, Total unimodularity and decomposition method for large-scale air traffic cell transmission model, *Transportation Research Part B* 53 (2013) 1–16.
- [66] V. Ramanujam, H. Balakrishnan, Increasing efficiency through flexibility in airport resource allocation, in: *Proceedings of the Conference on Decision and Control*, 2014.
- [67] T. Vossen, M. Ball, Slot trading opportunities in collaborative ground delay programs, *Transportation Science* 40 (1) (2006) 29–43.
- [68] H. Balakrishnan, Techniques for reallocating airport resources in adverse weather, in: *IEEE Conference on Decision and Control*, 2007.
- [69] S. L. Waslander, R. L. Raffard, C. J. Tomlin, Market-based air traffic flow control with competing airlines, *Journal of Guidance, Control, and Dynamics* 31 (1) (2008) 148–161.
- [70] S. L. Waslander, K. Roy, R. Johari, C. J. Tomlin, Lump-sum markets for air traffic control with competitive airlines, *Proceedings of the IEEE* 96 (2008) 2113–2130.
- [71] L. Li, J.-P. Clarke, A stochastic model of runway configuration planning, in: *AIAA Guidance, Navigation and Control Conference*, 2010.
- [72] C. A. Provan, S. C. Atkins, Optimization models for strategic runway configuration management under weather uncertainty, in: *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, 2010.
- [73] C. Weld, M. Duarte, R. Kincaid, A runway configuration management model with marginally decreasing transition capacities, *Advances in Operations Research*.
- [74] M. J. Frankovich, D. Bertsimas, A. R. Odoni, Optimal selection of airport runway configurations, *Operations Research* 59 (6) (2011) 1407–1419.
- [75] R. Oseguera-Lohr, N. Phojanamongkolkij, G. Lohr, J. W. Fenbert, Benefits assessment for tactical runway configuration management tool, in: *AIAA Aviation Technology, Integration, and Operations (ATIO) Conference*, 2013.
- [76] H. Hesselink, J. Nibourg, Probabilistic 2-day forecast of runway use: Efficient and safe runway allocation based on weather forecast, in: *Ninth USA/Europe Air Traffic Management Research and Development Seminar (ATM2011)*, 2011.
- [77] S. Houston, D. Murphy, Predicting runway configurations at airports, in: *Transportation Research Board (TRB) Annual Meeting*, 2012, paper number 12-3682.
- [78] M. Ben-Akiva, S. Lerman, *Discrete choice analysis: theory and application to travel demand*, MIT Press, 1985.
- [79] V. Ramanujam, H. Balakrishnan, Data-driven modeling of the airport configuration selection process, *IEEE Transactions on Human-Machine Systems* 45 (4) (2015) 490–499.
- [80] J. Avery, H. Balakrishnan, Predicting airport runway configuration, in: *Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2015)*, 2015.
- [81] M. Bierlaire, Biogeme: a free package for the estimation of discrete choice models, in: *3rd Swiss Transport Research Conference*, 2003.
- [82] A. Kannan, *A Distributed Feedback Control Model of the National Airspace System*, B.A. thesis, Harvard College (2009).

Hamsa Balakrishnan is an Associate Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology. She received her PhD in Aeronautics and Astronautics from Stanford University. Her research is in the design, analysis, and implementation of control and optimization algorithms for large-scale cyber-physical infrastructures, with an emphasis on air transportation systems. She received the US National Science Foundation’s CAREER Award in 2008, the Kevin Corker Award for Best Paper of the USA/Europe Air Traffic Management Seminar in 2011, the American Institute of Aeronautics and Astronautics’ Lawrence Sperry Award in 2012, and the American Automatic Control Council’s Donald P. Eckman Award in 2014.