

1
2
3
4
5
6
7
8
9 Applications of stochastic modeling in air traffic management: Methods,
10 challenges and opportunities for solving air traffic problems under
11 uncertainty*
12
13
14
15

16 Rob Shone[†]

17 Centre for Transport and Logistics (CENTRAL), Department of Management Science, Lancaster University Management
18 School, Lancaster University, Lancaster LA1 4YW, United Kingdom. Email: r.shone@lancaster.ac.uk.
19
20

21 Kevin Glazebrook

22 Centre for Transport and Logistics (CENTRAL), Department of Management Science, Lancaster University Management
23 School, Lancaster University, Lancaster LA1 4YW, United Kingdom. Email: k.glazebrook@lancaster.ac.uk.
24
25

26 Konstantinos G. Zografos

27 Centre for Transport and Logistics (CENTRAL), Department of Management Science, Lancaster University Management
28 School, Lancaster University, Lancaster LA1 4YW, United Kingdom. Email: k.zografos@lancaster.ac.uk.
29
30
31
32

33
34
35 **Abstract**
36
37

38 In this paper we provide a wide-ranging review of the literature on stochastic modeling applications within
39 aviation, with a particular focus on problems involving demand and capacity management and the miti-
40 gation of air traffic congestion. From an operations research perspective, the main techniques of interest
41 include analytical queueing theory, stochastic optimal control, robust optimization and stochastic integer
42 programming. Applications of these techniques include the prediction of operational delays at airports,
43 pre-tactical control of aircraft departure times, dynamic control and allocation of scarce airport resources
44 and various others. We provide a critical review of recent developments in the literature and identify
45 promising research opportunities for stochastic modelers within air traffic management.
46
47
48
49

50
51 *Keywords:* OR in airlines; Stochastic modeling; Stochastic optimization
52
53
54
55

56 *An earlier version of this work was included in the book of keynote papers for the OR60 annual conference at Lancaster
57 University, UK on September 11-13, 2018 (Shone et al. (2018)). Sections 2.2, 2.3, 2.4 and 3.1 include some material from the
58 conference paper. However, this manuscript as a whole represents a major expansion and enhancement of the aforementioned
59 conference contribution.
60

61 [†]Corresponding author
62
63
64
65

1 Introduction

Operations research (OR) methods for modeling and enhancing the efficiency of air traffic operations have made great advancements during recent decades. The range of applications in this area is vast; indeed, the management and optimization of an airport's performance requires consideration of a broad array of strategic, tactical and operational issues which can vary according to geographical and political circumstances (de Neufville and Odoni (2013), Zografos et al. (2013)). Furthermore, interactions between different airports (with respect to the propagation of flight delays, for example) imply that only limited insight can be gained by studying the operations of a single airport in isolation. The potential of OR methods to evaluate performance and identify improvement strategies in aviation settings has been demonstrated consistently over the last 60 years, although it could be argued that relatively few academic studies have sought to embrace the full range of complexities that might influence decision-making in practice.

The aims of this paper are to provide a broad literature review in order to demonstrate the fundamental role that stochastic modeling techniques have played in advancing aviation-related research, and also to discuss how these techniques can be applied to large-scale, dynamic, nonstationary optimization problems whose essential characteristics are not necessarily easy to identify or quantify. Many of the current problems of interest within air traffic management involve complicated sets of constraints, objectives and decision-making options, and the inclusion of uncertainty in such problems adds an extra layer of complexity due to the difficulties associated with selecting appropriate solution methods and modeling stochastic effects in a suitable way. Nevertheless, the advantages of being able to prescribe robust decision-making policies are too great to be overlooked, and the increasing availability of 'big data' is creating new opportunities for academic researchers to develop fine-tuned models of air transport operations. In order to emphasize the added value that stochastic modeling techniques can offer, we devote considerable attention in our review to research studies that have introduced elements of uncertainty to problems that were previously studied in *deterministic* settings, and suggest further opportunities for stochastic modeling to improve the quality of solutions or performance evaluations in similar environments.

Several existing research articles have provided high-level discussions of the different types of air traffic management problems in which OR methodologies (including both deterministic and stochastic modeling techniques) have the potential to make a significant impact. Many possible applications of stochastic modeling can be found in the area of *demand and capacity management*, which addresses the capability of an airport or airport network to efficiently handle the demands placed upon it by air traffic. The consequences of demand-capacity imbalances in the air transport system can be severe; indeed, in 2018, more than 19 million minutes of enroute delays were experienced by air passengers in Europe alone (International Air Transport Association (2019)). Research opportunities in this area, encompassing both strategic initiatives and tactical themes, have been discussed by Barnhart et al. (2012). More recently, Gillen et al. (2016) drew comparisons between administrative measures and economic incentives for aligning demand patterns with capacity limits, while Jacquillat and Odoni (2018) discussed the major interventions available to practitioners and used analytical insights to produce a roadmap for guiding policy and practice in the future.

In order to give our literature survey a clear focus, we define the scope of this paper according to a small number of specific, well-established mathematical themes that have risen to prominence in air traffic management and other OR application areas in recent decades. The survey will then examine the challenges, successes and opportunities associated with applying these OR themes to air traffic problems. The main themes of interest in this paper are as follows:

1. The use of stochastic queueing theory for modeling aircraft queues in capacitated settings in order to estimate operational delays at airports and in air traffic networks;
2. The use of robust and stochastic optimization to produce viable strategic, tactical or operational plans which hedge against the possible effects of uncertainty;
3. The use of stochastic optimal control methods to address sequential decision problems, including (for example) stochastic dynamic programming formulations for the allocation of scarce airport resources in response to the latest events and operating conditions on a particular day or season.

The above list might appear somewhat restrictive at first sight, but by focusing on research works which employ these techniques, we are able to examine a very broad section of the aviation literature which incorporates many prominent themes within demand and capacity management. Indeed, applications of the techniques listed above have been surprisingly wide-ranging. For example, simple queueing formulations were originally used in the 1950s to estimate the landing capacities of individual runways (Blumstein (1959)), but modern uses of queueing theory include the modeling of 'ripple effects' as delays propagate through networks of airports (Vaze and Barnhart (2012), Pyrgiotis et al. (2013)). Similarly, stochastic optimization methods for implementing ground delay programs were developed in the 1990s (Richetta and Odoni (1993)), but more recent applications have included large-scale air traffic flow management problems (Corolli et al. (2017), Jones et al. (2018)) and the tactical sequencing of aircraft take-offs and landings in order to optimize multi-criteria objective functions (Solak et al. (2018), Khassiba et al. (2020)).

As with any other application area in OR, the research landscape in air traffic management is defined not only by the development of new ideas for solving well-established problems but also by the emergence of new problems resulting from real-world changes in policy and practice. In recent years, the NextGen and SESAR projects (Joint Planning and Development Office (2010), European Commission (2018)) in the USA and Europe respectively have impacted the research agenda by promoting greater information sharing between airlines, air traffic controllers and other stakeholders. By exploiting the power of four-dimensional trajectory-based operations, individual flights should be able to achieve greater precision in meeting their pre-scheduled operation times (Hansen et al. (2009), Klooster et al. (2009), Dal Sasso et al. (2018, 2019)). These developments carry implications for the ways in which flight time uncertainty is modeled, as well as the formulations of tactical decision-making problems. However, these new technologies cannot eliminate the need for stochastic modeling; indeed, flight times and delays are still affected by a diverse range of factors including airframe-to-airframe variations in aerodynamic performance, variations in flight crew technique and limitations in wind prediction capability (Nikoleris and Hansen (2012)). The challenge for mathematical modelers is to recognize and adapt to technological advancements and policy changes in such a way that their studies remain relevant and (ideally) also achieve impact.

The remainder of this paper is organized according to different subapplication areas and problem types within air traffic management. Specifically, the sections are arranged as follows:

- Section 2 discusses queueing system formulations for air traffic;
- Section 3 discusses airport operations and capacity management, including airport scheduling mechanisms, dynamic allocation of airport resources and the control of airport surface operations;
- Section 4 discusses air traffic management and control problems, including the scheduling and sequencing of aircraft take-offs and landings and the dynamic control of flight trajectories in order to ensure safety and optimize on-time performance;
- Section 5 summarizes the main contributions of stochastic modeling discussed in the earlier sections and suggests several promising directions for future research.

Many of the topics discussed in our paper are inter-related and the organizational structure that we have outlined above is only one of several alternatives that could have been chosen, but we believe it is reasonably logical. From a practitioner's perspective, Section 2 is mainly relevant to performance evaluation, with queueing-theoretic methods for the estimation of operational delays under pre-determined schedules (at single airport and network level) presented and reviewed. Sections 3 and 4 place greater emphasis on optimization problems at the strategic, tactical and operational levels, with uncertainty of various kinds being an essential feature. Broadly speaking, Section 3 tends to focus primarily on the allocation of capacitated airport resources (e.g. take-off and landing slots) to airlines and aircraft, while Section 4 devotes greater attention to the detailed controls and adjustments that can be made to aircraft flight plans and trajectories in order to maximize efficiency whilst ensuring that safety standards are met. We also note that certain broad themes are examined from different perspectives in various parts of the paper; for example, airport capacity is an overarching concept which carries implications for queueing 'service time' distributions in Section 2, configuration of airport capacity envelopes in Section 3 and planned arrival acceptance rates under ground-holding programs in Section 4.

Table 1 provides a more detailed outline of Sections 2, 3 and 4 in our paper, with short descriptions included in order to summarize the principal topics of discussion within each subsection. It is intended that readers may use this table to see 'at-a-glance' which subsections are most relevant to the particular research questions that they may be interested in.

We have aimed to include some illustrative mathematical details (e.g. problem formulations) throughout the paper in order to provide a flavor of the OR techniques involved, although these details are not essential to our discussions and can be overlooked if desired.

In the online appendix to this paper we have provided a comprehensive set of summary tables in order to summarize some of the key attributes of aviation-related research articles cited in our survey. Three separate tables have been provided, corresponding to Sections 2, 3 and 4 in the main part of our paper. For example, the table corresponding to Section 2 summarizes specific attributes of articles related to queueing system formulations, including the type of queue(s) and the physical setting(s) being modeled. These tables are intended as an accompaniment to the discussions in our survey.

Table 1: An outline of the main discussion topics in each of the subsections in Sections 2, 3, 4

Sec.	Title and short description
2	Queueing system formulations of air traffic
2.1	Stochasticity and nonstationarity in air traffic queues <i>Fundamental considerations in air traffic queue modeling, including sources of uncertainty and nonstationarity, and (hence) the need for stochastic modeling with inclusion of time-dependent behavior</i>
2.2	Poisson demand processes <i>Applicability of Poisson processes for modeling the “joining” process in an air traffic queue; advantages and disadvantages of Poisson models; construction and validation of Poisson demand rate functions</i>
2.3	Pre-scheduled random demand (PSRD) queues <i>Applicability of ‘pre-scheduled random demand’ (PSRD) queues for modeling air traffic which operates according to schedules; relevant theoretical development; comparison with Poisson processes</i>
2.4	Modeling airport and runway capacity as a service process <i>Definition of airport/runway capacity and analytical methods for capacity estimation; geometric representation via capacity envelope; probability distributions for runway ‘service times’; queueing approximation algorithms</i>
2.5	Propagation of delays in an air traffic network <i>Challenges involved in modeling air traffic networks; iterative algorithms based on network decomposition and modeling of delays at individual nodes (airports); modeling congestion and delay propagation in enroute airspace</i>
3	Airport operations and capacity management
3.1	Slot allocation and the impact of slot limits <i>Administrative vs. market-based demand management strategies; MIP formulations for airport slot allocation; trade-off between schedule displacement and operational delays; incorporation of stochasticity</i>
3.2	Allocating runway capacity between arrivals and departures <i>Airport capacity utilization problem based on selection of ‘operating points’ on airport capacity envelope; use of stochastic queueing dynamics for measuring operational delays; solutions via DP and ADP methods</i>
3.3	Airport surface operations and departure control <i>Modeling of stochasticity in airport surface operations (e.g. push-back times, taxiing speeds) via data-driven methods and analytical queueing formulations; dynamic control of airport surface operations; gate assignment problems</i>
4	Air traffic management and control
4.1	Ground delay programs and air traffic flow management <i>Ground-holding problems, including static vs. dynamic and single airport vs. network; air traffic flow management problems; solution methods via stochastic integer programming, robust optimization, MDPs</i>
4.2	Runway scheduling problems <i>Aircraft sequencing problems, including static vs. dynamic, single operation type (e.g. arrivals only) vs. both types; ‘rolling horizon’ approaches; solutions via stochastic integer programming and robust optimization</i>
4.3	Aircraft conflict detection and resolution <i>Conflict detection and resolution over long, medium and short-term horizons; probabilistic approaches; geometric representation of conflict scenarios; dynamic control of aircraft speeds and conflict resolution maneuvers</i>

2 Queueing system formulations of air traffic

In reality, aircraft form queues when they are waiting to use the runways at airports. Air traffic queues can manifest in different ways; for example, in the case of aircraft waiting to land (arrivals), a ‘queue’ may refer to a holding stack above the terminal airspace, or a single-file stream of aircraft approaching a common runway, or even a backlog of aircraft being forced to delay entry to a congested air sector via ‘metering’ or ‘miles-in-tail’ interventions (Lulli and Odoni (2007)). Departure queues, resulting from airport surface congestion, can be observed on the taxiways and apron areas at busy airports worldwide (Simaiakis and Balakrishnan (2016)). Although aircraft queues can be defined in different ways according to the setting of interest, many of the corresponding mathematical models used by researchers share quite similar properties. Before discussing details of these, we observe that aircraft queueing models found in the literature commonly possess (at least) the following two characteristics:

- (i) *Stochasticity*: The times at which aircraft join the queues, and (usually) the lengths of service times in the queueing model, are subject to random variation;
- (ii) *Nonstationarity*: The queue dynamics (e.g. demand rates, service rates) vary with time due to their dependence on daily flight schedules and other controllable and uncontrollable factors, including active runway configurations and weather conditions.

We begin this section by discussing the motivation for including both stochasticity and nonstationarity in air traffic queueing formulations. This preliminary discussion can be found in Section 2.1. Subsequently, Section 2.2 discusses the continuing popularity of Poisson processes as a means for modeling air traffic demand processes; however, Section 2.3 counterbalances this by examining the recent trend for pre-scheduled random demand (PSRD) formulations to be used in order to overcome the limitations of a model with Poisson demand. Section 2.4 reviews the various ways in which researchers have used stochastic modeling to estimate runway capacity and airport capacity, and the assumptions that have been made regarding ‘service time’ distributions in queueing formulations. Finally, Section 2.5 discusses generalizations of the most popular airport queueing system formulations to network models.

Note on terminology: Throughout this section we use the term ‘arrival’ to refer to an aircraft arriving (i.e. landing) at an airport, and avoid using the term in a queueing-theoretic sense to refer to a new entrant in a queue. When discussing the processes or rates at which aircraft join queues (which may consist of arriving *or* departing aircraft), we refer to these as ‘demand’ or ‘joining’ processes/rates rather than arrival processes/rates in order to avoid possible confusion. Similarly, the term ‘departure’ is reserved for departing aircraft (take-offs).

2.1 Stochasticity and nonstationarity in air traffic queues

It is useful to consider the different sources and types of uncertainty that might influence the design of a queueing model in order to appreciate the benefits of accurate stochastic modeling in an air traffic context. Firstly, let us consider the process of new customers (or aircraft, in our context) entering a queue.

Passenger flights operate according to schedules, but - as in any transportation system - the adherence to scheduled operation times is not always precise. If the time at which an aircraft joins a particular queue (whether airborne or ground-based) differs from its scheduled entry time, and the deviation is not attributable to a particular control action (or multiple actions) within a mathematical model, then it is often treated as a stochastic effect. In the case where an aircraft is late to join a queue, we would normally refer to this as a delay. In the US, the Bureau of Transportation Statistics publishes on-time performance data which include the causes of flight delays (Bureau of Transportation Statistics (2020)). These causes are arranged into categories related to weather conditions, congestion effects in the airspace system, air carrier requirements (e.g. crew replacements), security issues and others. Within any of these categories, one can find a diverse range of factors which might be used to explain the stochasticity affecting the times at which new aircraft become 'present' (physically or otherwise) in a queue.

The nature of an aircraft 'service' process depends on the physical situation being modeled. If runway operations are being considered, then a 'service time' might refer to the amount of time between two consecutive aircraft clearing the runway, which depends on their runway occupancy times as well as the required separation distance involved. As we shall see later, service times are often treated as being less variable than queue-joining times in air traffic queueing models. Nevertheless, they are not entirely predictable. Indeed, runway occupancy times can be affected by aircraft weight, the condition of the runway surface, and (in the case of arrivals) the speed of touchdown and the pilot's choice of runway exit (Nikoleris and Hansen (2015), Meijers and Hansman (2019)). Separation requirements are also primarily determined by the weight classes of the aircraft involved; we discuss this further in Section 2.4.

It is clear that, although queue-joining and service times in aircraft queues can be influenced by many factors, not all of these factors should necessarily be treated as being *stochastic* in a formal sense. Indeed, by increasing the fidelity of a queueing model, one can often reduce the amount of stochasticity that remains. An obvious example to illustrate this point is the case of different aircraft weight classes. In a low-fidelity queueing model, customers (aircraft) may be treated as being homogeneous, in which case the effects of different separation requirements between different aircraft types may be witnessed as random variation in the model. On the other hand, if weight classes are accounted for in a high-fidelity queueing model via the inclusion of multiple customer classes, then much of this variation may be explained by the parameters for the different classes. Stochasticity can also be influenced by the time horizon of interest. For example, if one wishes to use a queueing model to predict the flight delays on a particular day three months in advance, then the weather conditions must be treated as unknown, and this carries implications for the specification of the queue entry and service time distributions (although there are obvious correlation effects to consider). On the other hand, if predictions are required for the next hour and the weather conditions are already known, then this particular source of random variation can be eliminated.

To some extent, therefore, it seems that 'randomness' can be controlled. Stochasticity can certainly be an artefact of the lack of detail in a queueing model, or the imprecision of the data used for its configuration. Experienced stochastic modelers are accustomed to the trade-off between tractability and fidelity in model design. While it may be tempting to include as much detail in a model as possible, this can hinder the efficient computation of useful summary performance measures. Indeed, omitting explanatory

variables at the expense of greater ‘stochastic variation’ is not necessarily a detrimental act if the resulting performance estimates are still fit for purpose. Furthermore, even the most high-fidelity model cannot realistically capture all of the human, mechanical and environmental influences that might affect aircraft queueing processes in practice. Thus, some level of uncertainty is always inherent. The queueing models discussed in this section feature obvious simplifications in many cases, but the results and insights offered by the associated research studies have provided sufficient value to justify their use.

It is a well-understood principle in queueing theory that higher levels of stochasticity tend to result in longer queues and longer waiting times. In the case of stationary queues, this can sometimes be verified directly using convenient formulae; for example, in the case of an $M/G/1$ queue, the Pollaczek-Khintchine formula shows that the average queue length is an increasing function of the service time variance (Gross and Harris (1998)). In the case of nonstationary queues, closed-form expressions for queue lengths and waiting times are not accessible in general, but it is possible to show experimentally that similar principles hold. For example, Hengsbach and Odoni (1975) (see also Kivestu (1976)) demonstrated that $M(t)/M(t)/k$ queues give consistently higher congestion estimates than $M(t)/D(t)/k$ queues when both models share the same time-varying expected service time values. Their results also suggest that, contrary to what one might expect, it is not necessary for demand rates to be particularly high in relation to service rates in order for $M(t)/M(t)/k$ results to differ significantly from the $M(t)/D(t)/k$ case.

The effects of nonstationarity on queueing performance measures are somewhat more difficult to analyze in a rigorous way. It has been shown in certain special cases that waiting times tend to decrease as the time-dependent demand rate becomes ‘more stationary’ (Ross (1978), Rolski (1981)), but in general it is very difficult to establish general principles which do not permit counter-examples (Heyman (1982)). However, the notion of nonstationarity having an adverse effect on system performance is corroborated by convincing evidence from numerical studies (Green et al. (1991), Schwarz et al. (2016)).

It is also important to note that, even if the underlying model parameters (such as demand rates and service rates) change only at fixed intervals of time, a queueing system may never come close to exhibiting ‘steady state’ behavior. Hence, piecewise stationary approximations may not be at all useful. Odoni and Roth (1983) examined the transient behavior of stationary, stochastic queueing systems. They considered the case of airport queues specifically and found that, during busy periods of the day, it can take many hours for steady state conditions to be reached - even if the demand rates and service rates do not vary at all during the relevant period. Thus, many situations of practical importance in air traffic applications require the use of transient (as opposed to stationary) queueing analyses.

Although most of the queueing formulations to be discussed in this section exhibit both stochasticity and nonstationarity, there have been several important research contributions that relied upon deterministic queue models and/or stationary queueing analyses, particularly in the classical literature. For example, Hubbard (1978) and Newell (1979) used cumulative diagrams based on deterministic, nonstationary queueing dynamics to model congestion-related delays at different times of day. Stationary, multi-class queueing models with stochastic dynamics were used by Rue and Rosenshine (1985) to search for optimal ‘balking points’ in order to set capacity thresholds for a single runway, and by Bauerle et al. (2007) (see also Bolender and Slater (2000)) to formulate routing problems in airports with multiple runways. Welch and

Lloyd (2001) used steady-state queueing performance measures derived from time-averaged demand rates to establish lower bounds for effective airport capacities. Marianov and Serra (2003) used properties of stationary $M/D/c$ queues to investigate the optimal locations of hubs in airport networks.

2.2 Poisson demand processes

As discussed in Section 2.1, many different factors can prevent air traffic operations from conforming to pre-determined schedules. Indeed, large deviations from scheduled operating times can occur as a result of flight cancellations, delays at ‘upstream’ airports, gate delays for departures, variability of flight times due to weather and winds, etc. (Pyrgiotis (2011), Belcastro et al. (2018)). One might therefore attempt to argue that, in practice, inter-joining times in air traffic queues are ‘sufficiently random’ to justify the use of nonhomogeneous Poisson models. In this type of model, the number of aircraft joining the queue between two points in time (say t_0 and t_1) is Poisson-distributed with a mean of $\int_{t_0}^{t_1} \lambda(t)dt$, where the demand rate function $\lambda(\cdot)$ is heavily dependent upon the schedule of operations. However, the use of a Poisson model requires some restrictive modeling assumptions. For example:

- (i) The variance of the number of queue entrants in any finite interval is equal to the mean;
- (ii) The numbers of queue entrants in two disjoint intervals are independent of each other.

It is clear that these are rather strong assumptions in the context of aircraft queues. If one considers the process of aircraft entering a terminal airspace, then it seems that a larger-than-expected number of entries during one hour should result in a smaller-than-expected number of entries during a different hour; similarly, a shorter-than-expected time gap between two consecutive entries is likely to cause a larger-than-expected gap elsewhere. Nevertheless, the apparent shortcomings of the Poisson model are arguably offset by its convenience and mathematical tractability. Dunlay and Horonjeff (1976) and Willemain et al. (2004) have used aircraft flight tracking data to provide empirical evidence in support of Poisson models. More recently, Wang et al. (2018) studied traffic flows in the US airspace system and found that a model with Poisson demand was able to match the ability of a more complicated Coxian queueing model to predict demand variations, with the additional benefits of greater tractability.

Poisson demand models for air traffic can be traced back to the classical literature. Galliher and Wheeler (1958) considered airport landings in the New York area and used the results of a case study to estimate demand rates over different time intervals during a typical day; thus, in their model, the demand rate $\lambda(t)$ has a piecewise constant structure. Koopman (1972) considered the case where arrivals and departures share a common runway, and proposed an extension whereby the Poisson demand rates for both operation types are not only time-dependent but also state-dependent. This results in Markovian state transition equations of the form

$$\begin{aligned} dP_{m,n}(t)/dt = & -(\lambda_{m,n}(t) + \eta_{m,n}(t) + \mu_{m,n}(t) + \nu_{m,n}(t))P_{m,n}(t) \\ & + \lambda_{m-1,n}(t)P_{m-1,n}(t) + \eta_{m,n-1}(t)P_{m,n-1}(t) \\ & + \mu_{m+1,n}(t)P_{m+1,n}(t) + \nu_{m,n+1}(t)P_{m,n+1}(t) \quad (m \geq 1, n \geq 1), \end{aligned}$$

where m and n are queue lengths for arrivals and departures, $P_{m,n}(t)$ is the corresponding state probability at time t , $\lambda_{m,n}(t)$ and $\eta_{m,n}(t)$ are queue entry rates and $\mu_{m,n}(t)$ and $\nu_{m,n}(t)$ are queue exit rates. This model allows for the possibility of ‘controlled’ demand rates, whereby the burden placed upon the system is reduced during peak congestion hours.

Hengsbach and Odoni (1975) extended Koopman’s approach to the case of multiple-runway airports, and claimed that results from a nonhomogeneous Poisson model were consistent with observed data from several major airports. In the last few decades, queues with nonhomogeneous Poisson demand processes have been used in a wide range of models and optimization problems based on single airports (Bookbinder (1986), Jung and Lee (1989), Daniel (1995), Hebert and Dietz (1997), Fan (2003), Mukherjee et al. (2005), Lovell et al. (2007), Stolletz (2008), Jacquillat and Odoni (2015a), Shone et al. (2019)) and also airport networks (Malone (1995), Long et al. (1999), Long and Hasan (2009), Pyrgiotis et al. (2013), Pyrgiotis and Odoni (2016)).

When using empirical data to design a nonhomogeneous Poisson demand process for a queueing model, the question arises as to how the time-dependent demand rate function $\lambda(t)$ should be constructed. The simple approach of using a piecewise constant function, as originally employed by Galliher and Wheeler (1958), continues to be used effectively in present-day applications (Jacquillat and Odoni (2015a), Jacquillat et al. (2017)), but one disadvantage of this approach is that ‘jump’ discontinuities occur at the endpoints of intervals when the demand rate changes. Hengsbach and Odoni (1975) avoided this problem by modeling the demand rate $\lambda(t)$ as a piecewise linear function, obtained by aggregating scheduled runway operations over each hour and then connecting the half-hour points using line segments, as shown in Figure 1. Bookbinder (1986) also used hourly data, but relied upon a three-point moving average method which automatically bridges the demand rates for different hours using sloping line segments. Clearly, many other interpolation methods are possible: however, a central feature of any Poisson demand model is that the variance of the entry count in any finite time interval corresponds to the mean, so (particularly during busy periods) the shape of the underlying demand rate function is likely to be obscured by random variation if one inspects a particular sample trajectory.

Many of the busiest airports around the world have multiple runways available, and these airports may elect to use certain runways for arrivals only, or for departures only. Queueing models for single runway operations may therefore restrict attention to either arrivals or departures, but ‘aggregated’ formulations which aim to model airports as single queueing systems must take into account both operation types. In many cases, an aircraft that lands at an airport will take off again (not necessarily from the same runway) within hours. This implies that the demand processes for arrivals and departures are not independent of each other, but in fact it is quite common in existing mathematical models for arrivals and departures to be treated as independent queues with their own time-varying Poisson demand rates. The assumption of independence is undoubtedly an oversimplification, but it may not be particularly harmful if one considers a large airport with separate runways being used for arrivals and departures (this system is referred to as ‘segregated operations’ and is used at London Heathrow, for example). Horonjeff and McKelvey (1994) (see also Grunewald (2016)) treated arrivals and departures as separate ‘job classes’ in a multi-class queue, but this approach is relatively uncommon in the literature.

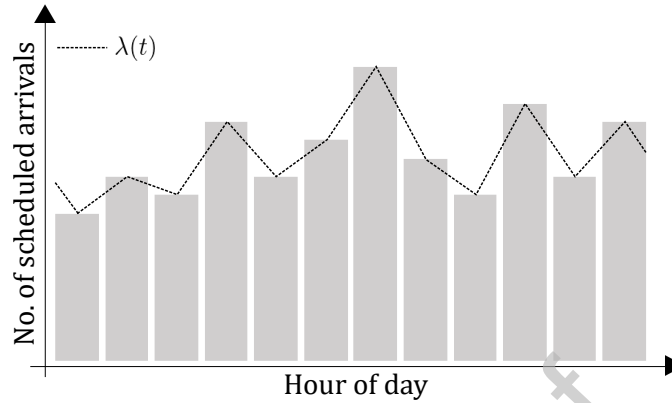


Figure 1: A piecewise linear demand rate function $\lambda(t)$, obtained by interpolating between half-hour points on a bar chart showing hourly demand (Shone et al. (2018))

It is interesting to note that the popularity of Poisson demand models for aircraft queues has led to their use for generating artificial data with which to demonstrate the performances of optimization algorithms. For example, in runway scheduling problems (discussed further in Section 4.2), the inter-joining times between consecutive aircraft entering the terminal airspace can be generated by sampling from exponential distributions. Some examples of this approach can be found in Beasley et al. (2004), Balakrishnan and Chandran (2010), Bennell et al. (2017) and references therein.

Nevertheless, with the emergence of new technologies that aim to mitigate unpredictable flight times (as discussed in the introduction), research trends appear to be shifting and researchers are beginning to show more interest in alternative models which allow greater control of inter-joining time variances. This is discussed further in the next subsection.

2.3 Pre-scheduled random demand (PSRD) queues

Queueing systems with pre-scheduled random demand (PSRD) are well-suited to applications in which customer queue entry times are dependent upon schedules. In PSRD queues (more commonly described as ‘pre-scheduled random arrivals’ or ‘PSRA’ queues in the literature), individual customers have their own pre-scheduled joining times but their actual joining times vary according to random earliness/lateness distributions; for example, deviations from scheduled times may be normally or exponentially distributed. Notably, PSRD queues are quite different from many of the classical models usually studied in queueing theory since inter-joining times are neither independent nor identically distributed. This implies that they do not belong to the family of demand processes referred to as ‘GI’ (General Independent) in Kendall’s classic notation (Kendall (1953)).

PSRD queues have been studied since the late 1950s (Winsten (1959), Mercer (1960)) but their application to aircraft queues is a relatively recent development. Ball et al. (2001) (see also Tu et al. (2008)) conceived the idea of applying random perturbations to scheduled departure times in a Ground Delay Program (GDP) application. Subsequently, Guadagni et al. (2011) attempted to study PSRD queues more

rigorously. They used the expression

$$t_i = \frac{i}{\lambda} + \xi_i, \quad i \in \mathbb{Z},$$

to represent the actual entry time of the i^{th} customer to join the queue, where $1/\lambda$ is the expected inter-joining time between two consecutive customers and the random variables ξ_i are independent and identically distributed with variance σ^2 . This model can easily be modified to allow for cancellations and unexpected queue entries ('pop-ups'), making it more suitable for air traffic applications. Guadagni et al. (2011) observed that, although this process is known to converge weakly to a Poisson process as σ increases, its behavior can be very different from the Poisson case if σ is large but finite. In particular, if one defines N_1 and N_2 as the numbers of queue entries in time intervals $[t, t+u]$ and $[t+u, t+2u]$ (respectively), then

$$\text{Cov}(N_1, N_2) = - \sum_i p_i^{(\sigma)}(t, t+u) p_i^{(\sigma)}(t+u, t+2u),$$

where $p_i^{(\sigma)}(t, t+u)$ is the probability that the i^{th} customer will join during interval $(t, t+u)$ given that $\text{Var}(\xi_i) = \sigma^2$. Hence, PSRD queues exhibit negative autocorrelation, in the sense that time intervals which experience fewer queue entries than expected are likely to be followed by time intervals with more entries than expected (and vice versa).

Applications of PSRD queues have become increasingly common in the last decade. Nikoleris and Hansen (2012) considered normally distributed deviations from scheduled times and used the Clark approximation (Clark (1961)) to estimate expected queueing delays. Caccavale et al. (2014) used a PSRD model to study inbound traffic at Heathrow Airport and argued that Poisson processes are a poor model for arrivals at a busy airport since, in practice, the arrivals stream is successively rearranged according to air traffic control rules. Gwiggner and Nagaoka (2014) compared a PSRD model with a Poisson model using a case study based on Japanese air traffic and found that the two models exhibited similar behavior in systems with moderate congestion, but deviated from each other during high congestion. Lancia and Lulli (2020) studied the arrivals process at eight major European airports and found that a PSRD model with nonparametric, data-driven delay distributions provided a better fit for the observed data than a Poisson model.

In recent years, PSRD queues have also been incorporated within optimization problems. Furstenu et al. (2015) studied a runway scheduling problem¹ under uncertainty and used a PSRD process, calibrated using German airport data, to model arrival delay times. Shone et al. (2019) considered Poisson and PSRD processes as two different cases for the queue-joining processes in a sequential decision problem formulated as a Markov decision process, with airport queues as the primary application. In the PSRD case, deviations from scheduled times were assumed to be exponentially distributed in order to improve mathematical tractability.

The main advantages of using PSRD formulations to model aircraft queues are:

- (i) They allow variances of deviations from scheduled times to be controlled independently of expected values, so that the process of aircraft joining a queue can be made arbitrarily close to a deterministic

¹Runway scheduling problems are discussed further in Section 4.2.

process, even when the queue is very busy;

- (ii) They capture the negative autocorrelation effect that one would expect in queueing systems where customer entry times are dependent upon a pre-determined schedule;
- (iii) Individual aircraft can be given their own earliness/lateness distributions, allowing greater modeling flexibility and differentiation based on aircraft types (for example).

On the other hand, an obvious disadvantage of PSRD formulations is that they are generally less mathematically tractable than queues with Poisson demand, although in certain special cases (e.g. the case of exponentially-distributed lateness) one can still derive time-differential equations similar to the Chapman-Kolmogorov equations for Markovian queues.

2.4 Modeling airport and runway capacity as a service process

An airport's capacity can be defined as the expected number of runway movements (either arrivals or departures) that can be operated per unit time under conditions of continuous demand (de Neufville and Odoni (2013)). This leads to a natural comparison between the capacity of an airport (or a single runway) and a service rate in queueing theory. For strategic purposes, it is important for airport controllers and policymakers to estimate an airport's capacity accurately, since efficient runway operations (without the airport's resources being either under-utilized or overburdened) rely upon the attainment of a suitable balance between demand and capacity. However, the situation is complicated by the fact that airport throughput rates are affected by a number of time-varying factors, including weather conditions, runway configurations and different air traffic mixtures. Hence, capacity estimation at airports (with resulting implications for service rates in queueing models) is worthy of academic study in its own right.

Blumstein (1959) used a stochastic model, featuring randomly-distributed approach speeds and separation requirements between different aircraft types based on air traffic control (ATC) standards, to estimate the landing capacity of a single runway. This work was generalized by Hockaday and Kanafani (1974), who derived expressions for runway capacity under three different modes of operation (arrivals only, departures only and mixed operations), with time separations between consecutive landings and runway occupancy times modeled as being normally distributed and type-dependent.

A key principle that emerged from these early contributions was the importance of taking into account different possible fleet mixes and sequencing strategies. In practice, aircraft are categorized into different 'weight classes', and the minimum permitted time separation between two consecutive runway movements depends on these weight classes (and also the type of operation: either 'arrival' or 'departure'). Heavier aircraft invariably generate more 'wake turbulence', and hence the time separation required is greater if a heavy aircraft is followed by a small aircraft (Newell (1979)). Therefore, in airport capacity calculations, one must take into account the relative expected frequencies of different 'weight pairs' (e.g. heavy-small, heavy-heavy etc.) and use these to calculate average time separations between consecutive movements.

Gilbo (1993) developed the idea of a *runway capacity curve* (referred to by subsequent authors as a 'capacity envelope'), as shown in Figure 2. This curve represents the departure capacity of an airport as

a convex, nonincreasing function of the arrival capacity. The shape of the curve depends on various time-varying factors, including weather conditions and the runway configuration in use. However, the essential principle is that each point on the capacity envelope represents a feasible pair of capacity values for arrivals and departures during the time period for which the envelope applies. Various authors have provided detailed descriptions of how airport capacity envelopes can be constructed using analytical methods (Lee et al. (1997), Stamatopoulos et al. (2004)) and empirical methods (Ramanujam and Balakrishnan (2009), Simaiakis (2013), Ju et al. (2015)) and these capacity envelopes have been incorporated into various types of optimization problems, which are discussed further in Sections 3 and 4.




Figure 2: A piecewise linear capacity envelope under two possible cases for the prevailing weather conditions: VMC (good weather) and IMC (bad weather)

The capacity envelope representation allows service rates for arrivals and departures to be chosen as points in the two-dimensional plane that satisfy the linear constraints shown in Figure 2. However, the use of a queueing model for airport traffic also requires selection of a suitable service time distribution. The seminal paper of Galliher and Wheeler (1958) (which considered arrivals only) used an $M(t)/D(t)/c(t)$ model, with a nonhomogeneous Poisson process for demand and deterministic service times, but since then randomly-distributed service times have been widely adopted. Koopman (1972) suggested that the queueing dynamics of an airport with s runways (modeled as independent servers) could be bounded by the characteristics of the $M(t)/D(t)/s$ and $M(t)/M(t)/s$ queueing systems. The former system can be regarded as a ‘best-case’ scenario, since queueing delays are shorter in the case of predictable service times, while the latter system - with exponentially-distributed service times - is a ‘worst-case’ scenario. In these formulations, numerical solution of the Chapman-Kolmogorov equations (assuming a finite queue capacity) can be used to estimate time-dependent queue length probability distributions.

Kivestu (1976) proposed an $M(t)/E_k(t)/s$ queueing model for aircraft queues, in which the service time distribution is Erlang with k exponentially-distributed service phases. This approach is closely related to that of Koopman (1972), since the cases $k = 1$ and $k = \infty$ represent exponential and deterministic service times respectively. However, Kivestu also introduced a fast, practical numerical approximation for the time-

dependent queue length probabilities in an $M(t)/E_k(t)/1$ queue, which became known as the DELAYS algorithm. This method involves the construction of a set of carefully-designed *pseudo-completion* epochs, based on the expected service completion times of customers. To be specific, one defines the n^{th} epoch t_n by

$$t_n = \frac{k+1}{k} \sum_{i=1}^n \bar{x}_i,$$

where \bar{x}_i is the expected service duration of the i^{th} customer (which depends on their queue-joining time and the time-varying service rate). Notably, the pseudo-completion epochs become more closely-spaced as k increases, which ensures that the queue length and waiting time estimates given by the DELAYS algorithm tend to decrease as the service time variation is reduced. If the probability of j queue entries during interval $[t_n, t_{n+1})$ is approximated by

$$\alpha_{n+1}(j) = \frac{\left(\frac{\lambda(t_n)}{\mu(t_n)}\right)^j \exp\left(-\frac{\lambda(t_n)}{\mu(t_n)}\right)}{j!},$$

(where $\lambda(t)$ and $\mu(t)$ are time-dependent demand and service rates respectively), then one can use recursive equations of the form

$$P_j(t_{n+1}) = P_0(t_n)\alpha_{n+1}(j) + \sum_{i=1}^{j+1} P_i(t_n)\alpha_{n+1}(j-i+1) \quad (j \geq 0)$$

to estimate the queue length probability distributions at epochs t_n for $n \geq 0$. The DELAYS algorithm - as well as the $M(t)/E_k(t)/1$ model itself - has become very popular, and has been used for estimating air traffic delays in a variety of settings (Abundo (1990), Malone (1995), Fan and Odoni (2002), Stamatopoulos et al. (2004), Mukherjee et al. (2005), Churchill et al. (2008), Hansen et al. (2009), Pyrgiotis and Odoni (2016)). The algorithm offers the advantage of generating full (estimated) queue length probability distributions at different points in time. This is useful because, in reality, aviation practitioners are not only interested in the *expected* queueing delays at different times of day; they are also interested in tail-based performance measures such as the probabilities of queueing delays exceeding certain thresholds (e.g. 15 minutes).

An advantage of using Erlang-distributed service times is that they can accurately approximate many other non-Markovian or empirical distributions, due to the 'richness' of the Erlang family of distributions (Gupta (2010)). However, the $M(t)/E_k(t)/1$ model for aircraft queues can be regarded as somewhat macroscopic, since it does not explicitly take into account fleet mixes and separation requirements between different aircraft types; instead, it assumes that such considerations can be implicitly accounted for by adjusting the phase parameter k (which controls the service time variance). Models that explicitly consider runway occupancy times for different classes of aircraft have been proposed by various authors. Hockaday and Kanafani (1974) and Stamatopoulos et al. (2004) modeled these using normal distributions, while Jeddi et al. (2006) suggested beta distributions and Nikoleris and Hansen (2015) used Gumbel random variables.

The assumption of a single server (as in the $M(t)/E_k(t)/1$ model, for example) is surprisingly common

in the literature, even when modeling operations at a multiple-runway airport. One possible explanation for this is that, even when an airport has multiple runways, there is usually some inter-dependence between them, which implies that it is inappropriate to model them as independent servers (Jacquillat (2012)). For example, runways may intersect each other - or even if they do not, they may be too closely-spaced to allow independent operations, since the effects of wake turbulence may create ‘diagonal separation requirements’ between aircraft on adjacent runways (Stamatopoulos et al. (2004)). Nevertheless, the single-server assumption is a simplification that has arguably been over-used in the literature.

The DELAYS algorithm has proven to be an effective means of approximating the dynamics of an $M(t)/E_k(t)/1$ queue, and it is noted here that alternative methods are also available. Malone (1995) proposed the ‘State Probability Vector Approximation’ (SPVA), which is somewhat similar to DELAYS but has the additional advantage of being applicable to more general service time distributions. Stolletz (2008) implemented the ‘stationary backlog-carryover’ (SBC) method, which is based on modifying queue-joining rates according to ‘blocking’ probabilities in a queue with no waiting room and can also be applied to queues with general inter-joining and service time distributions. Lovell et al. (2013) described the use of diffusion approximation methods for modeling delays at a single airport, and explained how this approach might be extended to a network scenario. Shone et al. (2019) discussed the idea of using a ‘server always busy’ approximation to model the queueing dynamics of an $M(t)/E_k(t)/1$ or $PSRD/E_k(t)/1$ queue during periods of time in which demand exceeds capacity (which can often be the case at congested airports). More generally, methods for modeling time-dependent queues are discussed in Green et al. (2007), Defraeye and Van Nieuwenhuyse (2016) and Schwarz et al. (2016).

Finally, we note that the notion of “capacity” has been used in quite a context-specific way in this subsection in order to suit the purposes of the queueing formulations (in particular, the service time distributions) under discussion. In Sections 3 and 4 we address a broad range of decision problems that fall into the general area of *airport capacity management*, in which the deployment of limited airport resources in response to (or anticipation of) time-varying patterns of demand is of primary interest.

2.5 Propagation of delays in an air traffic network

Models of queueing delays in airport networks have made somewhat limited progress in comparison to models of delays at individual airports. This is partly because it is very difficult to validate the outputs of network queueing models using real data. As observed by several authors (see Churchill et al. (2008, 2010), Arikan et al. (2013) and references therein), occurrences at a single airport may not be easily translatable into observed effects at another, and this hinders the application of traditional queueing network models. For example, an aircraft that takes off an hour late from its origin airport may not necessarily land an hour late at its destination; instead, it may be able to expedite its arrival in order to make up for lost time (Bratu and Barnhart (2006), Kohl et al. (2007), Selim Akturk et al. (2014)). Nevertheless, the propagation of delays around an airport network is an important phenomenon to study and this subsection discusses the progress that has been made in this area using stochastic queueing models.

Arguably the most promising approach to have emerged from the literature in recent years involves iterating between two distinct stages. In the first stage, queues at individual airports are modeled using

static schedules, as if each airport is operating in isolation. In the second stage, schedules at individual airports are updated based on the queueing delays computed in the first stage. Peterson et al. (1995) provided early inspiration for this approach by considering a network configured in a hub-and-spoke fashion, with a single hub airport at its center. Their approach uses deterministic queue dynamics, but allows the capacities of airports to vary according to a semi-Markov process. Following the calculation of expected delays at individual airports, the *updated* scheduled arrival time for a particular flight f at airport n is given by

$$\tau_{nf} + (d_f + \alpha \mathbb{E}[W_{t-1}^n] + (1 - \alpha) \mathbb{E}[W_t^n] - s_{nf})^+,$$

where τ_{nf} is the *original* scheduled arrival time, d_f is the total of the cumulative delays incurred by flight f over previous flight legs, $\mathbb{E}[W_t^n]$ is the mean waiting time for an aircraft arriving at airport n at the end of discretized time period t , s_{nf} is the amount of ‘slack time’ incorporated into the schedule for the journey to airport n (which may be used to nullify some of the delays) and $\alpha \in [0, 1]$ is a weight that depends on the precise timing of arrival. This update rule can be used in an iterative procedure, with the expected delays $\mathbb{E}[W_t^n]$ for airports $n = 1, 2, \dots, N$ re-calculated at each discrete time period as the scheduled flight arrival times are updated.

Subsequent models have used the $M(t)/E_k(t)/1$ formulation to model delays at individual airports (Long et al. (1999), Long and Hasan (2009)). The ‘Approximate Network Delays’ (AND) model, first conceptualized in Malone (1995) but developed further in Pyrgiotis et al. (2013), is also based on $M(t)/E_k(t)/1$ queues and iterates between a queueing engine (the DELAYS algorithm) and a delay propagation algorithm that accounts for the ‘slack times’ typically built into airport schedules; see Figure 3.

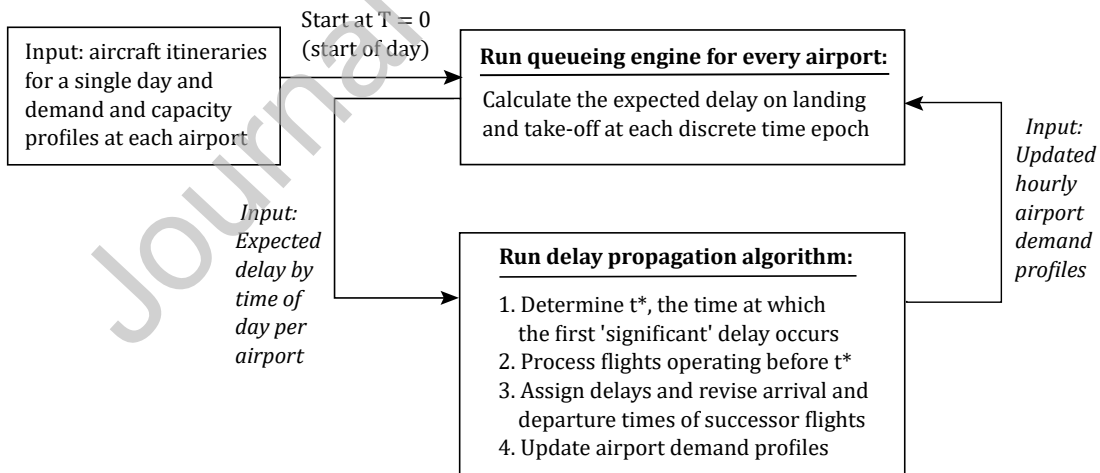


Figure 3: The Approximate Network Delays (AND) algorithm, adapted from Pyrgiotis et al. (2013)

The AND algorithm has been used effectively in applications including the assessment of demand management strategies at congested airports (Vaze and Barnhart (2012), Pyrgiotis and Odoni (2016)) and analyses of the effects of ‘local disturbances’, such as weather events or air traffic controller strikes

(Baspinar et al. (2016)). However, a number of authors have pointed out weaknesses of the AND algorithm that could be addressed in future work. Lovell et al. (2013) observed that the AND queueing engine (based on $M(t)/E_k(t)/1$ queues) assumes that the input process to each ‘downstream’ airport is a Poisson process, even though the output process from an ‘upstream’ airport might be very different from the Poisson case. Thus, it would be preferable to have a mechanism of controlling the variances of inter-joining times independently of their means. Pyrgiotis et al. (2013) remarked that AND does not allow for interventions that airlines might make in response to congestion - which might include the cancellations of severely delayed flights or the swapping of delayed aircraft. Additionally, there may be causes of delay that are unrelated to queueing dynamics, such as mechanical or crew scheduling issues.

The references given above are mainly concerned with the modeling of delays at airports. From the perspective of air traffic flow management (ATFM), queueing network models can also be used to model stochastic demand patterns and congestion effects in different airspace sectors, and there is a considerable body of research in this area (see Sridhar et al. (2008) for a useful review). Shortle et al. (2003) pointed out the limitations of using analytical methods which assume Poisson demands at individual nodes and suggested an efficient simulation model with fewer binding assumptions. Tandale et al. (2008) proposed an $M/M/c$ Jackson queueing network to model the interactions between airspace network flows. Wan et al. (2011) (see also Tien et al. (2011), Zhou et al. (2011), Taylor and Wanke (2013) and references therein) created a highly configurable queueing network model of the US National Airspace System (NAS) which allows for stochastic processes (e.g. Poisson processes) to model the demand at individual nodes, and presented the results of simulation experiments. They also considered probabilistic weather variations. Optimizing the routing decisions in large networks is a formidable challenge, but these models have strong potential to be used for decision support (Taylor et al. (2012), Wanke et al. (2012)).

Models of network delay propagation may be used in order to assess the impact on global air traffic delays of a particular configuration of airport schedules, or (at a more tactical level) a particular assignment of ground-holding and/or air-holding delays as part of an air traffic flow management policy. The airport slot allocation problem is described in Section 3.1, while air traffic flow management is discussed in greater detail in Section 4.1.

3 Airport operations and capacity management

The queueing system formulations discussed in Section 2 underpin much of the work to be discussed in the next two sections. In this section we focus on *airport operations and capacity management*, which includes themes related to flight scheduling at a strategic level and the control of scarce runway capacity and airport surface operations at a tactical level. Many of the research studies that we discuss in this section are concerned with the formulation and solution of optimization problems, with objectives generally related to the efficiency of airport capacity utilization.

Specifically, this section is organized as follows:

- Section 3.1 discusses the strategic problem of airport slot coordination, and the inherent trade-off that exists between satisfying airlines’ scheduling requests and maintaining acceptable limits on

operational delays (taking into account stochastic influences on delays);

- Section 3.2 discusses the dynamic allocation of runway capacity in order to generate time-varying ‘service rates’ for queues of arrivals and departures at a single airport;
- Section 3.3 discusses the control of airport surface operations, which can be accomplished via the use of dynamic push-back release policies and gate assignment strategies (for example).

Many of the optimization problems considered in this section are related to *demand and capacity management* (discussed in our introduction). We note that the term ‘demand’ was used in Section 2 for quite a specific, queueing-theoretic purpose, i.e. to represent the process of aircraft entering a queue. In this section, we use ‘demand’ in a rather more general way; for example, it might refer to the number of requests received from airlines to use a particular airport time slot for take-off or landing. We aim to avoid any possible confusion by providing enough context in our discussions to ensure that important terms such as ‘demand’ and ‘capacity’ have unambiguous meanings.

3.1 Slot allocation and the impact of slot limits

The busiest airports outside the US fall into the category of slot-controlled (level 3) airports, which means that airlines intending to use these airports for take-offs or landings must submit requests for time slots (typically 15 minutes long) during which they have permission to use the runways and other airport infrastructure. Although the US does not implement slot controls in the same manner, a small number of its airports are subject to scheduling limits which restrict the number of hourly runway movements (Zografos et al. (2017), Jacquillat and Odoni (2018)).

Since slot allocation is usually carried out with a broad set of objectives in mind (including the need to design schedules that satisfy airlines’ requirements as equitably as possible), the resulting schedules do not always insure effectively against the danger of severe operational (queueing) delays. For example, if too many flights are allocated to a small set of consecutive time slots, the consequences for airport congestion levels may be catastrophic. Thus, there is a need for demand management strategies to ensure that congestion mitigation is included as part of the slot allocation procedure.

A useful survey of demand management strategies that have been implemented in the US, Europe and other parts of the world is provided by Fan and Odoni (2002). These strategies can generally be divided into two categories: administrative and market-based. Administrative strategies involve setting ‘caps’ on the numbers of runway operations that can take place at an airport in a single time period, or a number of consecutive time periods. These ‘caps’ may apply to arrivals, departures or both, and are usually referred to in the aviation community as declared capacities (Zografos et al. (2017)). On the other hand, market-based strategies are based on using economic measures such as congestion pricing and slot auctions to relieve congestion during peak periods (Andreatta and Odoni (2003), Fan (2003), Pels and Verhoef (2004), Mukherjee et al. (2005), Ball et al. (2006, 2020), Andreatta and Lulli (2009), Pellegrini et al. (2012)). A number of authors have directly compared administrative and market-based strategies using analyses and/or case studies (Brueckner (2009), Basso and Zhang (2010), Czerny (2010), Gillen et al. (2016)).

Market-based strategies for mitigating airport congestion appear to have a lot of potential, but the relevant econometric methods of analysis do not easily fit within the scope of this paper as defined in Section 1; furthermore, these methods have yet to gain popularity in practice. This subsection will therefore focus on administrative strategies (slot controls). Let us consider the situation at a slot-controlled airport. *Schedule displacement* (or sometimes *schedule delay*) is a term used to represent the lack of conformity between the confirmed flight schedule for a single airport (or network) and the set of slot requests originally made by airlines. For example, if one uses r_f to denote the requested operation time for an individual flight f (either an arrival or a departure) and t_f to denote the time given to flight f in the final schedule, then the schedule displacement can be measured simply by

$$\sum_{f \in \mathcal{F}} |t_f - r_f|,$$

where \mathcal{F} is the set of flights to be scheduled on an individual day or a scheduling season. The real-world problem of allocating flights to scheduling slots in such a way that various constraints (based on declared capacities at airports, ‘turnaround time’ requirements for individual aircraft, fairness considerations, etc.) are satisfied has been modeled by researchers as an optimization problem which is typically solved using mixed integer programming methods (Zografos et al. (2012), Pellegrini et al. (2017), Ribeiro et al. (2018, 2019), Zografos and Jiang (2019), Fairbrother et al. (2020), Fairbrother and Zografos (2020)). This subsection discusses how stochastic modeling considerations can be incorporated into this process.

Various authors (Barnhart et al. (2012), Swaroop et al. (2012), Zografos et al. (2012)) have commented on the inherent trade-off that exists between schedule displacement and operational (queueing) delays, as illustrated by Figure 4. At slot-controlled airports, certain time slots tend to be more sought-after by airlines than others. As a result, flight schedules that conform closely to airline requests are likely to result in large ‘peaks’ in demand at certain times of day. These schedules incur only a small amount of schedule displacement, since the requests from airlines are largely satisfied; however severe operational delays are likely to be caused by the peaks in demand. Conversely, operational delays can be reduced by smoothing (or ‘flattening out’) the schedule to avoid such peaks, but this generally involves displacing flights to a greater extent from the times requested by airlines and thereby results in more schedule displacement.

Two-stage stochastic optimization has been proposed by Corolli et al. (2014) and Wang and Jacquillat (2020) as a suitable method for incorporating stochasticity into slot allocation problems. Corolli et al. (2014) considered a network-level problem in which the first-stage objective function has the form

$$\text{Min} \sum_m \sum_t d_m c_m^t x_m^t + W \cdot \mathbb{E}[f(\mathbf{x}, \tilde{\omega})],$$

where d_m is the number of days of the scheduling season on which a specific runway movement m is requested, x_m^t (a decision variable) indicates whether or not movement m is scheduled for a particular time slot t and c_m^t is the corresponding schedule displacement cost. The cost of operational delays in the resulting slot allocation is $f(\mathbf{x}, \tilde{\omega})$ (weighted by the parameter W), which is calculated in the second-stage




Figure 4: The trade-off between schedule displacement and operational (queueing) delays (Shone et al. (2018))

problem by

$$\text{Min} \sum_a q_a \sum_d \sum_t (y_{adt}^\omega + z_{adt}^\omega).$$

Here, q_a is the cost of delaying a movement at airport a , and y_{adt}^ω and z_{adt}^ω are decision variables that represent the numbers of delayed departures and arrivals (respectively) on day d and time instant t after the realization of the random variable ω , which defines a set of airport capacities. The variables y_{adt}^ω and z_{adt}^ω must satisfy lower-bounding constraints determined by the mismatch between demand and capacity under random scenario ω . As with all stochastic optimization formulations, it is critical that the uncertainty set from which scenarios are drawn is specified carefully according to knowledge of the real-world situation; otherwise, there is no guarantee that the solutions obtained will be more effective in practice than those that would be obtained using a simple deterministic formulation.

Wang and Jacquillat (2020) have also adopted a stochastic programming framework in order to optimize schedule interventions under weather-related uncertainty. Their objective function is comparable to that of Corolli et al. (2014), since it includes separate components for schedule displacement and expected delay costs. After airport capacities are realized, operational decisions are made in the form of ground delays to impose on individual flights (we discuss ground holding problems further in Section 4.1). Their model is applied to the full US network, and this requires the development of original decomposition methods in order to ensure computational tractability. Their paper makes significant modeling and methodological

contributions, although it also leaves scope for future research; for example, additional types of air traffic flow management decisions (such as rerouting or cancellations) have not yet been incorporated.

Churchill et al. (2013) used stochastic optimization methods to determine the optimal numbers of slots to make available at different times during a day, taking into account different probabilistic airport capacity profiles based on historical weather patterns. Their approach is notable because it addresses the optimal specification of declared capacities, rather than treating these as being ‘inputs’ to the optimization model. Pyrgiotis and Odoni (2016) used a case study to demonstrate how, given a set of slot scheduling constraints, one can mitigate operational (queueing) delays by ‘smoothing’ demand rates whilst also complying with these constraints. Their study used the AND model (see Section 2.5) to model aircraft queues stochastically.

Recently, a very interesting trend to have emerged in the literature has been the incorporation of uncertainty based on stochastic queueing dynamics into slot allocation problems. As described in Jacquillat and Odoni (2018), stochastic queueing systems exhibit highly non-linear relationships between demand rates (or service rates) and expected queueing delays. This makes it very challenging to incorporate queueing dynamics into integer programming formulations. Furthermore, demand rates at airports (or airport networks) are mainly controlled at the *strategic* level via slot coordination, whereas service rates are (to a certain extent) within the realm of tactical control, since airports can allocate runway capacity dynamically between arrivals and departures according to the latest observed congestion levels. Attempts to optimize demand-capacity relationships at airports should therefore take into account the various different kinds of interventions that are possible at different stages of a planning horizon, as well as the specialized stochastic modeling techniques required to predict operational delays accurately.

Jacquillat and Odoni (2015a) have developed a new framework for airport slot allocation which is considerably more ambitious in its approach to the modeling of stochastic operational delays than previous work. Their approach relies upon a *collinearity assumption*, which essentially states that if one particular schedule is preferable to another with respect to on-time performance when queueing dynamics are modeled *deterministically*, then the same should be true when the queueing dynamics are modeled *stochastically*. This assumption cannot be shown to be correct in all cases, but empirical evidence suggests that it is often correct. Thus, even though deterministic queue dynamics will always underestimate the delays that would occur under stochastic conditions (Hansen et al. (2009)), one may be able to employ them in order to simplify the search for an optimal schedule with respect to on-time performance. Jacquillat and Odoni (2015a) used an iterative approach, in which the first stage uses a mixed integer programming formulation to optimize on-time performance under constraints based on maximum permitted scheduling displacements to individual flights and deterministic queue dynamics. Subsequently, tactical interventions (e.g. service rate adjustments) are used to optimize on-time performance under the optimal schedule obtained from the first stage, with queues modeled *stochastically*. The stochastic modeling of queues for arrivals and departures is based on the $M(t)/E_k(t)/1$ formulation, discussed earlier in Section 2.4. If the resulting performance is not satisfactory, then the first stage can be revisited and the scheduling displacement constraints relaxed in order to enable the optimization model to find a superior schedule with respect to queueing performance. This model has also been extended in order to incorporate inter-airline equity objectives (Jacquillat and Vaze (2018)).

The approach of Jacquillat and Odoni (2015a) appears to be very promising, although it is somewhat dependent on the validity of the collinearity assumption, which is difficult to verify in any formal sense. Future research opportunities might involve the integration of stochastic queueing dynamics directly into an optimization model using a nonlinear programming approach, or the careful derivation of ‘surrogate’ queueing performance metrics that might fit within a linear programming formulation. Whether or not such approaches would be computationally tractable in any realistically-sized problem, however, is unclear. If an iterative approach (involving a deterministic model and a stochastic model, as in Jacquillat and Odoni (2015a)) is found to be the only practical way forward in such problems, then there is certainly scope for attempting to use fine-grain simulations or alternative stochastic models (based on PSRD queues, for example) in order to examine a proposed schedule’s on-time performance more rigorously.

In the general stochastic programming literature, there is a well-known performance measure called the *value of the stochastic solution* (VSS) which measures the advantage that one is able to gain by incorporating knowledge of probability distributions into an optimization procedure (Birge and Louveaux (2011)). Specifically, the VSS is the difference (in percentage terms, for example) between the objective function value given by a stochastic optimization procedure and the performance of an ‘expected value’ (EV) solution under stochastic conditions, where the EV solution is obtained by simply setting all stochastic parameters to their expected values and then using a deterministic optimization method. The VSS can be investigated in the context of slot allocation, although (as in any stochastic programming model) it is highly dependent on the parameter values and distributions involved.

In the model of Corolli et al. (2014) discussed earlier, the parameter W effectively controls the importance of operational delays relative to schedule displacement. Since the operational delays depend on probabilistic capacity scenarios (while the schedule displacement has no stochastic behavior), it is natural to expect both the amount of schedule displacement and the VSS to increase with W . In their experiments, Corolli et al. (2014) found that $W = 4$ was sufficient to yield VSS values greater than 50% in some problem instances. Wang and Jacquillat (2020) found that the VSS in their problem instances varied between 4% and 23.1%, with the relative weight of operational costs (as opposed to scheduling costs) again being a key influence.

Given its problem-dependent nature, the VSS is perhaps less interesting to discuss than the qualitative differences between solutions prescribed by stochastic and deterministic optimization methods. Assuming a weighted objective function of the kind used by Corolli et al. (2014) and Wang and Jacquillat (2020), we conjecture that deterministic optimizers will tend to prescribe smaller amounts of schedule displacement than stochastic optimizers, since they rely on ‘averaged’ or nominal estimates of airport capacities (or other stochastic parameters) which neglect the very high operational costs that one might incur under the ‘worst-case’ scenarios. In other words, the expected value of operational delays under a particular schedule will tend to include a disproportionately large contribution from the worst-case scenarios. A stochastic optimizer should be able to perform well by taking these kinds of tail-based effects into account. Our conjecture is supported by the results in Corolli et al. (2014), and we also note that the iterative algorithm in Jacquillat and Odoni (2015a) implicitly relies upon these kinds of principles by gradually relaxing the schedule displacement constraints until the stochastic performance becomes acceptable. We discuss the

VSS and related issues again later, in other problem contexts.

Before concluding this subsection, we would like to suggest one further possible future research direction. It is well-known that airlines often use ‘schedule padding’ to insure against the effects of congestion-related delays by allowing more time for individual flight legs to be completed and thereby improving their on-time performance. This practice appears to be particularly common in the US, where delays can be especially unpredictable due to the lack of slot controls at airports (Ball et al. (2010a), Odoni et al. (2011)). However, schedule padding incurs costs of its own, since long buffer times can prevent airlines from making efficient use of their aircraft, crew, infrastructure and other scarce resources; moreover, there may be a cost associated with lower consumer demand (Skaltsas (2011), Yimnga and Gorjidoz (2019)). Optimizing the amount of ‘padding’ to include in scheduled flight legs requires econometric analyses, but stochastic modeling can also play a substantial role since the risks involved are related to probabilities of being able to meet on-time performance targets. To the best of our knowledge, this type of problem has received relatively little attention in the literature thus far.

3.2 Allocating runway capacity between arrivals and departures

The use of an *airport capacity envelope* for representing the interdependence between maximum achievable throughput rates for arrivals and departures has been discussed in Section 2.4. This approach was first proposed by Gilbo (1993) (see also Gilbo (1997)), who also formulated a sequential decision problem in which a period of operations is divided into discrete time slots (e.g. 15 minutes long) and the decision-maker is able to control the arrival and departure ‘capacities’ in each time slot by selecting an appropriate point on the capacity envelope. In Gilbo (1993), the decision-maker’s objective is to minimize a function of the form

$$\sum_{i=1}^N \gamma_i [\alpha_i X_{i+1}^k + (1 - \alpha_i) Y_{i+1}^k],$$

where N is the number of time slots, X_i and Y_i are queue lengths for arrivals and departures respectively at the beginning of slot i and $\alpha_i \in [0, 1]$ and $\gamma_i \geq 0$ are weight parameters. In realistic problems, the α_i values will tend to be greater than 0.5 due to the greater costs associated with airborne delays (as opposed to ground delays), while γ_i might be decreasing with i due to the greater uncertainty associated with operating conditions in more distant time slots (indeed, γ_i may be likened to a ‘discount factor’ in discrete-time Markov decision processes; see Puterman (2005)). If $k = 1$ then one obtains a linear function of the queue lengths, but various factors might motivate a difference choice of k in practice; for example, Jacquillat et al. (2017) used a model in which the expected total delay scales quadratically with the number of queueing aircraft.

The problem formulated by Gilbo (1993) can be treated as a dynamic problem, with the choice of service rates for time slots $j \geq i$ adjusted at the beginning of each slot $i \in \{1, 2, \dots, N\}$ according to the latest available information on expected demand rates, meteorological conditions, etc. However, a significant limitation of Gilbo’s model is that the cost minimization is carried out under the assumption of *deterministic* queue dynamics. A number of other research contributions have also considered capacity utilization problems in a deterministic setting (Hall (1999), Gilbo and Howard (2000), Dell’Olmo and Lulli

(2003), Weld et al. (2010), Bertsimas et al. (2011a)). The formulation of stochastic, dynamic optimization problems for airport capacity utilization is a relatively new development in the air transport literature. Jacquillat et al. (2017) considered a discrete-time Markov decision process (MDP) in which the system state at the beginning of time slot i is represented by a 5-tuple

$$(a_i, d_i, RC_i, wc_i, ws_i),$$

where a_i and d_i are observed queue lengths for arrivals and departures respectively, RC_i is the runway configuration in use, wc_i represents the weather state (which affects the shape of the capacity envelope) and ws_i represents the wind state (which affects the set of runway configurations available). Notably, the state variables RC_i , wc_i and ws_i are restricted to small, discrete sets, which helps to ensure that the state space is not too large to facilitate a rigorous dynamic programming (DP) approach. After observing the system state, the decision-maker chooses service rates for arrivals and departures (with the set of possible choices being discretized) and also has the option of changing to a different runway configuration. From a stochastic modeling perspective, the most important innovation of the model used by Jacquillat et al. (2017) is the incorporation of stochastic queueing dynamics based on independent $M(t)/E_k(t)/1$ queues for arrivals and departures. Furthermore, transition probabilities are calculated via the rigorous numerical solution of Chapman-Kolmogorov equations, without an approximation method (e.g. the DELAYS algorithm described in Section 2.4) being required. Additional stochasticity is incorporated by using Markov chains to model the random evolution of weather and wind conditions.

The model proposed by Jacquillat et al. (2017) (see also Jacquillat and Odoni (2015a,b)) may be regarded as the first serious attempt to formulate the problem of optimizing runway capacity utilization as a discrete-time MDP via the use of stochastic queueing dynamics. By restricting the sizes of the state and action spaces and using discretization where necessary, the authors were able to use a conventional DP approach to find optimal dynamic policies in a case study based on JFK Airport in New York (see also Zambon (2018) for a European case study). They also investigated the benefit (in terms of congestion cost savings) of modeling the queue dynamics stochastically rather than deterministically, and found that this ranged between 5% and 20%. The smallest cost savings were found to occur in problem instances where runway configuration changes could be made without any cost - thereby allowing the deterministic model to react to unanticipated system state transitions by exploiting a costless control mechanism.

Shone et al. (2019) considered a somewhat similar model to that of Jacquillat et al. (2017), but removed the requirement for the action space (consisting of feasible service rates) to be discretized. By making use of an approximation for the dynamics of an $M(t)/E_k(t)/1$ queue that becomes increasingly accurate as the demand-to-capacity ratio increases, they were able to propose an approximate dynamic programming (ADP) approach that bypasses the need for transition probabilities to be computed explicitly. They also considered PSRD processes (see Section 2.3) as an alternative to the Poisson demand model, but did not explicitly consider variable weather conditions or include runway configuration changes as a control mechanism.

It appears that there is considerable scope for the use of ADP methods to prescribe strong decision-making policies in tactical decision-making problems with complicated state and action spaces. In reality,

the information available to decision-makers in such problems is extensive; for example, the latest positions and estimated times of arrival (ETAs) of aircraft enroute to an airport are known to air traffic controllers, and these are subject to dynamic uncertainty. Furthermore, in-air separation requirements between aircraft are dependent upon their weight classes, and therefore airport capacities (typically translated into service rates in queueing models) are strongly dependent upon the mixtures of different types of aircraft using the runways. It is not computationally feasible to use exhaustive DP methods to solve problems with vast state and action spaces; however, many different types of ADP methods (based on artificial neural networks, for example) might be used to find *approximately* optimal policies in such problems (see Bertsekas and Tsitsiklis (1996), Sutton and Barto (1998), Powell (2007) for background information). These methods appear to have a lot of unexplored potential in airport capacity utilization problems.

3.3 Airport surface operations and departure control

Most of the queueing models discussed in this paper so far have been related to runway operations (take-offs and landings), without explicit consideration of the fine-grain operations involved in maneuvering aircraft so that they are ready to join a runway queue. However, bottlenecks can also occur away from the runways; for example, departing aircraft might experience delays caused by congestion on the airport taxiways. This section examines how stochastic modeling has been used with respect to airport ground operations, with a particular focus on aircraft departure processes.

As noted in Section 2, $M(t)/E_k(t)/1$ queues have been used extensively to model queues of arrivals and departures at airports. Although empirical studies have been provided to support the assumption of Poisson demand processes for arrivals (Dunlay and Horonjeff (1976), Willemain et al. (2004)), we are not aware of any similar attempt to validate the Poisson model for airport departures. The main factors that affect aircraft departure times have been incorporated within simulation studies (Shumsky (1995), Clarke et al. (2007)) and statistical models for prediction (Idris et al. (2002), Carr (2004)). These factors include gate departure delays (which can be caused by passenger delays, crew scheduling issues, mechanical failures etc.), interaction effects between different runways (which are particularly relevant if runways intersect each other, or if they are parallel but in close vicinity of each other) and adverse weather conditions. Recently, Badrinath et al. (2020) have provided motivation for stochastic modeling approaches by demonstrating the significant impact of demand-related uncertainty on airport surface operations.

Pujet et al. (1999) (see also Andersson et al. (2000)) used a data-driven queueing model for airport departures, with stochasticity introduced via the use of Gaussian distributions to model push-back durations, taxiing speeds and other factors. Simaiakis and Balakrishnan (2009) also used data-calibrated Gaussian random variables to model ‘unimpeded’ taxi-out times of aircraft, with actual taxi-out times obtained by including the effects of congestion. Subsequently, Simaiakis and Balakrishnan (2016) extended this work by developing and testing a model that predicts runway schedules and take-off times in response to a given aircraft push-back schedule. Notably, their approach makes use of a $D(t)/E_k(t)/1$ queueing model. Aircraft travel times between departure gates and runways are estimated via a separate procedure in order to generate an expected runway schedule, which then provides the deterministic, time-varying demand rates for the stochastic queueing model. The use of Erlang-distributed service times in their model is supported

by an earlier empirical study (Simaiakis and Balakrishnan (2013)) which demonstrates the advantages of such an approach. Other researchers have used alternative methods for predicting aircraft taxi-out and departure delays under uncertainty; Balakrishna et al. (2010) employed reinforcement learning (RL) methods, while Ravizza et al. (2013) have used regression-based analyses.

The research described above is related to the prediction of aircraft departure delays via stochastic queueing or data-driven methods. Naturally, these kinds of modeling approaches can also be used to formulate decision problems. Burgain et al. (2009) used an MDP approach to optimize the control of the push-back and taxiing processes under different levels of information regarding aircraft positions. Simaiakis et al. (2014) also considered a dynamic control problem in which decisions are made regarding time-dependent push-back rates. These push-back rates then act as inputs to an $M(t)/E_k/1$ runway queueing model. They considered system states of the form (R_t, Q_t) , where R_t is the number of aircraft taxiing to the runway at the start of discrete time epoch t and Q_t is the length of the runway queue (measured in terms of Erlang service phases). Since all of the taxiing aircraft are assumed to have reached the runway by the start of the next epoch, the Bellman equations can be written in the simplified form

$$J^*(r, q) = \text{Min}_{\lambda} \left\{ c(r, q) + \alpha \sum_{j=0}^{kC} \mathbb{P}_{(r,q) \rightarrow (\lambda, j)} J^*(\lambda, j) \right\},$$

where λ is the number of aircraft that push back during period t , $c(\cdot)$ is a single-step cost function, J^* is the optimal cost-to-go function, α is a discount factor and C is the finite queue capacity. Optimal push-back policies are then obtained using DP policy iteration methods.

Badrinath and Balakrishnan (2017) studied optimal control policies in a system of two queues in tandem, with the first queue representing congestion in an airport ramp or apron area and the second representing runway congestion. Although the queueing dynamics of their model are governed by simple differential equations (with optimal push-back policies obtained by solving a deterministic nonlinear program), simulation experiments are used to test the performances of the resulting push-back policies in a stochastic environment. McFarlane and Balakrishnan (2016) considered a similar dual-queue model and also investigated the effects of using different time discretizations for decision-making purposes. Lian et al. (2019) have demonstrated the benefits of dynamic push-back control policies using data obtained from Beijing International Airport. Their study includes the use of an iterative algorithm to optimize the choice of threshold K in an $M/M/1/K$ model for an airport taxiway queue. Chen and Solak (2020) considered a problem in which departing aircraft can be held either at a designated metering area or at the gates, and sought to optimize traffic flows at different surface locations under operational uncertainty.

Another type of surface management problem that one might consider is a gate assignment problem, in which flights must be assigned to departure gates under various ‘strict’ constraints (e.g. the need to avoid two flights being assigned to the same gate concurrently) and ‘softer’ constraints (e.g. the assignment of gates in such a way that flights operated by the same airline are located in the same physical area of the airport). Typical objectives might include minimizing the number of un-gated flights (i.e. flights assigned to the apron area), minimizing the towing operations required, or minimizing total passenger walking distance. A useful survey of such problems is provided by Bouras et al. (2014).

Although gate assignments are commonly affected by unforeseen disruptions, it appears that only limited attention has been given to stochastic versions of this problem. Some authors have used robust optimization, with a certain amount of ‘buffer time’ included in gate departure schedules in order to absorb stochastic delays (Mangoubi and Mathaisel (1985), Hassounah and Steuart (1993), Yan and Chang (1998)). Alternative robust optimization methods are proposed by Lim et al. (2005), Dorndorf et al. (2007) and Yan and Tang (2007). In a similar vein, Narciso and Piera (2015) (see also Yan et al. (2002)) have used simulation experiments to evaluate the robustness of different gate assignment policies. Seker and Noyan (2012) developed a stochastic optimization approach, with uncertainty related to flight arrival and departure times (treated as model inputs). Aoun and El Afia (2014) proposed an MDP formulation of a stochastic gate assignment problem, in which transition probabilities are based on potential conflicts (caused by operational delays) between flights assigned to the same gate.

Optimal control of airport surface operations is a problem that, if desired, can be formulated at a very microscopic level - with consideration of the availability of ground vehicles, apron stands, etc. Like other problems discussed in this paper, it is also a problem which (ideally) should not be treated in isolation, as there are obvious implications for aircraft departure times and other relevant performance indicators. In Section 4 we discuss various topics that carry implications for airport surface operations, including ground delay programs and the sequencing and scheduling of runway operations.

4 Air traffic management and control

In this section we discuss topics related to the detailed control of air traffic operations. The related optimization problems are based on many different types of decision-making options available to air traffic controllers and airport coordinators, including the routes taken by individual aircraft, the assignment of take-offs and landings to airport runways and the adjustments to flight trajectories needed in order to minimize conflicts. Thus, many of the research studies discussed in this section can be differentiated from those considered in Section 3 by their somewhat more ‘fine-grain’ consideration of controls exercised on individual flights; however, it should be noted that there are many synergies between the topics discussed in different sections of our survey, and indeed several of the papers that we shall discuss in this section also consider themes relevant to topics discussed in previous sections.

Specifically, this section is organized as follows:

- Section 4.1 discusses the use of ground delay programs and (more generally) the important role of air traffic flow management in optimizing the efficiency of operations;
- Section 4.2 discusses the sequencing and scheduling of airport runway operations in order to meet (for example) on-time performance or fuel consumption objectives;
- Section 4.3 discusses the planning and real-time control of aircraft flight trajectories in order to prevent or resolve mid-air conflicts whilst minimizing disruptions to flight plans.

4.1 Ground delay programs and air traffic flow management

In previous sections we have discussed the implications of demand-capacity imbalances, with particular attention to the situation at airports with limited physical infrastructure. The slot allocation problems discussed in Section 3.1 were strategic problems in which decisions were required long before any operational uncertainty (e.g. flight cancellations, weather events) could be realized. However, interventions can also be made on a day of operations in order to manage demand rates dynamically at congested airports and air sectors. This subsection briefly discusses the literature on *air traffic flow management* (ATFM) problems, and describes how researchers have modeled uncertainty in such problems.

ATFM problems were first conceptualized by the seminal work of Odoni (1987), which described a stochastic, dynamic problem in which flows of traffic must be managed according to available capacity in an airspace system. More recently, Vossen et al. (2011) have provided a broad overview of current practices in ATFM and outlined the relevant mathematical models and optimization problems. From a methodological perspective, the techniques of interest in this area are mainly related to stochastic programming, with uncertainty introduced by allowing airport and air sector capacities to depend probabilistically on different weather scenarios. Deterministic and stochastic formulations of ATFM problems are currently attracting a lot of research interest, and in this subsection we aim to offer a simple introduction to the topic by focusing on the stochastic *Ground Holding Problem* (GHP) - which is arguably the simplest type of ATFM problem - in order to give a flavor of the modeling and solution methods involved. Subsequently, the use of stochastic optimization in more general ATFM problems is briefly discussed.

The GHP is based on the principle of replacing expensive airborne delays with ground-holding delays. To be more specific, if a particular flight is expected to arrive at its destination airport during a period of heavy congestion, then it may be safer and more cost-effective (with respect to fuel consumption, etc.) to delay its departure from the origin airport in order to ensure that arrival occurs during a less-congested period. Stochasticity is introduced to the problem by considering weather effects. In poor weather conditions, an airport's capacity can be reduced significantly due to the longer aircraft separation times required. Odoni et al. (2011) studied empirical data from major European and US airports and found that the reductions in hourly capacity due to adverse weather were mainly in the range 10%-15%.

Andreatta and Romanin-Jacur (1987) were among the first authors to use a stochastic optimization approach. In later work, Terrab and Odoni (1993) formulated the GHP in such a way that the objective is to minimize the function

$$\sum_{i=1}^N Cg_i(X_i) + \sum_{q \in \mathcal{Q}} \left[p_q \sum_{i=1}^N Ca_i(X_i, Y_i^q) \right],$$

where p_q is the probability of the capacity scenario $q \in \mathcal{Q}$ occurring, N is the number of flights to be (potentially) delayed on the ground, $Cg_i(x)$ is the 'ground-holding' cost of delaying flight $i \in \{1, 2, \dots, N\}$ for x time periods before take-off, and $Ca_i(x, y)$ is the cost of delaying flight i in the air by y time periods if it has already been delayed by x periods on the ground. A solution (or policy) is represented by a vector (X_1, \dots, X_N) , whereas the vector (Y_1^q, \dots, Y_N^q) is dependent on airborne delays under capacity scenario q . In

order to represent the real-world problem accurately, it must be the case that $C a_i(x, y) + C g_i(x) > C g_i(x+y)$ for any x and y . Terrab and Odoni (1993) found that a stochastic dynamic programming approach for optimizing the vector (X_1, \dots, X_N) was not computationally feasible, and instead presented heuristic methods.

The formulation described above assumes that ground-holding delays can be assigned to flights individually. However, like other problems in air traffic management, the GHP has been modified over the years in order to reflect technological and policymaking changes in the air transport system. In 1998, radical changes in ground delay programs were implemented for all US airports as a result of the collaborative decision-making (CDM) initiative, which allowed for much greater interaction between airlines and the Federal Aviation Administration (FAA) (Ball et al. (2000), Barnhart et al. (2003)). Consequently, research efforts shifted towards considering GHP formulations in which control is exercised on groups of flights. Ball et al. (2003) proposed a new formulation of the GHP to comply with the CDM paradigm. Subsequently, Kotnyek and Richetta (2006) showed that the formulation of Ball et al. (2003) was closely related to an earlier formulation used by Richetta and Odoni (1993), in which the objective is to minimize

$$\sum_{q \in \mathcal{Q}} p_q \left[\sum_{i=1}^T \sum_{j=i+1}^{T+1} C_g(j-i) X_{qij} + c_a \sum_{i=1}^T W_{qi} \right],$$

where $C_g(i)$ denotes the ‘ground-holding’ cost of delaying an aircraft for i periods before take-off, c_a is the marginal cost of air delay per aircraft (assumed constant), W_{qi} is the number of aircraft unable to land during period i under capacity scenario q due to congestion at the destination airport, and X_{qij} (a decision variable) is the number of aircraft originally scheduled to arrive during period i but re-scheduled to arrive during period j under scenario q . Richetta and Odoni (1993) also considered an extension to multiple aircraft classes and solved this problem using integer programming techniques. Kotnyek and Richetta (2006) showed that solutions to the linear programming relaxation of this model possess the integrality and equity properties required to be valid under the CDM paradigm, provided that certain conditions on the ground-hold cost functions are met.

The formulations used by Terrab and Odoni (1993) and Richetta and Odoni (1993) (see also Vranas et al. (1994b)) are valid only for *static* versions of the GHP, which do not allow solutions to be updated dynamically during the day. Solution methods for dynamic versions of the problem, including the development of strong-performing heuristics, were first developed by Richetta and Odoni (1994), Vranas et al. (1994a) and Richetta (1995). Mukherjee and Hansen (2007) (see also Liu et al. (2008)) generalized these earlier models by allowing ground-holding decisions to be revised at different decision-making stages in response to evolving information about airport capacities (modeled via scenario trees). Recently, Estes and Ball (2020) have further strengthened the formulations of Mukherjee and Hansen (2007) and Ball et al. (2003) by improving their scalability and also allowing the option of flight diversions.

Ball et al. (2010b) investigated the idea of using a priority scheme that improves GHP performance by giving preferential treatment to long-haul flights, and proposed heuristic methods for ameliorating the inequity introduced by such a scheme. Glover and Ball (2013) built upon this work by using a two-stage

stochastic optimization model with a multi-objective function based on trade-offs between efficiency and inter-flight equity. Jones et al. (2015) showed that improvements in GHP performance can be achieved if arrival times (as opposed to departure times) are controlled by means of enroute speed adjustments. Many of the aforementioned papers rely upon the use of probabilistic capacity profiles based on weather conditions, and various algorithmic and simulation-based methods for generating these have been proposed; see Inniss and Ball (2004), Buxi and Hansen (2011), Clarke et al. (2013) and references therein.

In more general ATFM problems, control actions can include the imposition of ground or airborne delays and the cancellation or rerouting of flights after uncertainty has been realized, while uncertainty is usually related to the capacities of airports and air sectors and is often modeled via scenario trees in order to allow events to unfold dynamically. Other sources of uncertainty (e.g. with respect to flight arrival times) can also be incorporated (Jones et al. (2018)). A useful classification and overview of ATFM problems has been provided by Agustin et al. (2010). Typical solution methods are based on stochastic programming, including the use of two-stage and multi-stage formulations with recourse; see Alonso et al. (2000), Clarke et al. (2009), Ganji et al. (2009), Mukherjee and Hansen (2009), Andreatta et al. (2011), Agustin et al. (2012b), Balakrishnan and Chandran (2014), Chang et al. (2016), Corolli et al. (2017) and references therein. Some authors have also considered chance-constrained programming (Clare and Richards (2012), Chen et al. (2017), Yang (2018)), while robust optimization approaches have received a relatively limited amount of attention thus far (Gupta and Bertsimas (2011), Saraf et al. (2012, 2014)). The problem of dynamically optimizing aircraft flight paths subject to weather-based uncertainty can be formulated as a Markov decision process (Nilim et al. (2001), Nilim and El Ghaoui (2004)).

For the purposes of context, we also note here that some of the most important contributions in ATFM research considered deterministic settings (Bertsimas and Stock Patterson (1998), Lulli and Odoni (2007), Bertsimas et al. (2011b), Agustin et al. (2012a)) and these have provided important foundations for many of the stochastic formulations that have followed.

The use of stochastic programming formulations, with probabilistic scenarios typically related to airport capacities and air sector capacities, makes it easy (and often instructive) to investigate the VSS (defined in Section 3.1) in GHPs and ATFM problems. The VSS varies a lot with different model formulations and, even within a particular formulation, can be highly sensitive to certain model parameters or conditions. Chang et al. (2016) found that the VSS ranged between 6% and 30% depending on the volume of flights included in their single-sector ATFM problem, with greater volumes generally yielding higher values. Corolli et al. (2017) found that the VSS was closely related to the timing of decisions for a particular subset of flights in their model. Specifically, if complete flight plans (including routing, ground-holding and air-holding decisions) are required in the first stage of their problem, then the VSS can be as high as 14%, but if the holding decisions are not required until the second stage (after the realization of capacity-related uncertainty), then the VSS is usually much smaller (e.g. less than 1%).

If deterministic optimization values are used in a stochastic ATFM problem (with capacity profiles averaged over all possible scenarios, for example), then there may be a risk of obtaining infeasible solutions. Alonso et al. (2000) and Clare and Richards (2012) have noted that the solutions obtained using such methods can violate capacity constraints. At the other end of the scale, however, deterministic optimiza-

tion methods can also yield overly conservative solutions. This has been noted by Clarke et al. (2009), who suggested that under-utilization of airspace capacity could occur as a result of aircraft being held on the ground for too long or sent on excessively long routes under a deterministic model, without due consideration of the possibility that shorter routes might become available as a result of weather improvements (for example). Similarly, Andreatta et al. (2011) found that stochastic solutions tend to allow higher numbers of departures in earlier time periods, whereas deterministic models implement more ground delays in these early periods. In this way, the stochastic solutions are able to ‘hedge’ against uncertainty by ensuring that if scenarios without capacity reduction actually occur, then capacity is not wasted. The VSS was found to be about 6% on average in their computational experiments.

The results of Corolli et al. (2017) support the view that stochastic optimization methods offer greater benefits if decisions are subject to higher levels of uncertainty. Indeed, several of the classical GHP papers compared static and dynamic decision-making problems, with the latter allowing decisions to be made with the latest available information. Richetta and Odoni (1993) found that the VSS was about 7% in the static version of their GHP model, but much greater cost improvements were possible by allowing decisions to be made dynamically (Richetta and Odoni (1994), Richetta (1995)). Indeed, dynamic solutions are able to update ground-holding decisions based on the latest capacity forecasts and obtain much lower air delay costs. Mukherjee and Hansen (2009) considered a static GHP formulation but allowed rerouting decisions to be made dynamically. They found that a dynamic rerouting model could yield cost improvements of 10%-15% compared to a static model which forces earlier rerouting decisions. Notably, the dynamic model is able to achieve substantial savings in ground delay costs (cf. Andreatta et al. (2011)) by releasing flights towards blocked entry fixes and subsequently rerouting them if necessary.

In an ATFM context, it seems reasonable to surmise that the benefits of stochastic optimization tend to diminish if decisions are allowed to be made more frequently. However, most practical situations do not allow decisions to be made infinitely often, and indeed it is not advisable to give air traffic controllers overwhelming workloads. Thus, decisions must inevitably be made under some level of uncertainty in such environments. In Section 4.2 we consider runway scheduling problems, in which the advantages of stochastic optimization methods can be evaluated using similar principles. In Section 4.3 we discuss the subfield of aircraft conflict detection and resolution, which also has close links to ATFM.

4.2 Runway scheduling problems

Many of the optimization problems discussed earlier in this paper are related to the control of air traffic flows at a somewhat macroscopic level (via the specification of demand or service rates in a stochastic queueing model, for example). This is certainly the case for the capacity utilization problems discussed in Section 3.2, and also for some of the surface operations models in Section 3.3. This subsection devotes attention to *runway scheduling*, in which the precise ordering of aircraft using a runway system needs to be determined. We note that runway scheduling problems have also been referred to as *aircraft sequencing* problems (ASPs) in the literature, and also *aircraft landing* problems (ALPs) in cases where only arrivals are considered. However, the term “runway scheduling” now seems to be becoming more popular as it reflects the broader range of decision options that might be included; for example, decisions might include

the allocation of landings or take-offs to specific runways in addition to aircraft sequencing.

Runway scheduling problems (RSPs) typically possess certain characteristics which differentiate them from other types of capacity utilization problems. The essence of such problems is to determine an optimal sequence of ‘runway movements’; these movements could consist of arrivals only, departures only or a mixture of both, depending on the problem description. Usually, the individual flights that need to be ‘sequenced’ possess their own specific attributes; for example, a particular flight may belong to a certain weight class, with its own preferred landing (or take-off) time that depends on the flight carrier’s operational requirements. After passing a certain temporal or spatial threshold, an aircraft’s position in the runway sequence may be considered ‘frozen’, i.e. no longer adjustable; for example, Figure 5 depicts a situation where aircraft of different weight classes arrive from different directions at an ‘entry fix’ near the terminal airspace. After passing the entry fix, they proceed along a common glide path in first-come-first-serve order; hence, sequencing decisions apply only to aircraft that have not yet progressed beyond the entry fix.

By taking into account the time separation requirements between different aircraft weight classes, an RSP decision-maker must find a runway sequence that optimizes an objective function based on either a single criterion or multiple criteria. Typical performance criteria are based on adherence to scheduled operation times, fuel consumption costs (measured by the ‘holding times’ incurred by aircraft before they are added to the runway sequence) and the total amount of time required for all runway operations to be completed (Bennell et al. (2017)).

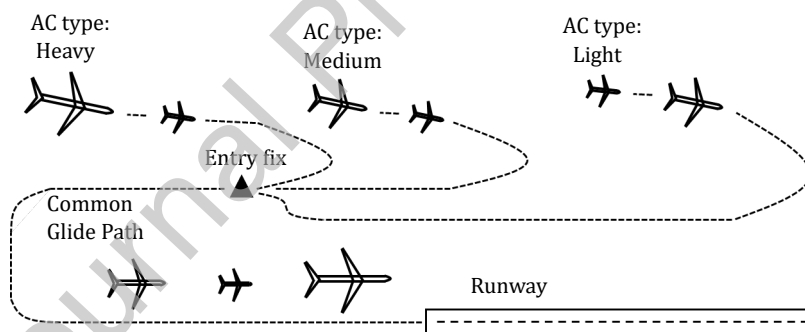


Figure 5: A runway scheduling problem with three different aircraft types, adapted from Hu and Chen (2005)

The literature on *deterministic* RSPs is well-developed, and we will provide only a few references here in order to provide context for the stochastic modeling approaches that have emerged in recent years. Psaraftis (1978) was among the first authors to consider a ‘static’ RSP in which all relevant information is known in advance, so that there is no need to update sequences dynamically. Dear (1976) formulated a dynamic version of the problem and also introduced the widely-adopted concept of ‘constrained position shifting’ (CPS), which imposes constraints on the amount by which a particular flight is allowed to deviate from its position in a first-come-first-served (FCFS) sequence. It is quite common for dynamic RSPs to be solved using a ‘rolling horizon’ approach, in which the runway sequence is updated at fixed time intervals (e.g. five minutes) in order to allow new aircraft to enter consideration as and when they become ‘ready’ for take-off or landing; see Cieselski and Scerri (1997), Beasley et al. (2004), Hu and Chen (2005), Moser

and Hendtlass (2007), Murca and Muller (2015), Bennell et al. (2017) for examples.

The ‘rolling horizon’ approach may be regarded as a means of accounting for the effects of uncertainty in an RSP, since it allows problem parameters (e.g. aircraft ETAs) to be updated at each time increment according to the latest available information. However, one might also attempt to model stochastic effects explicitly, rather than employing deterministic solution methods in a stochastic environment. Several computational studies have provided motivation for this type of approach. Stamatopoulos et al. (2004) demonstrated the benefits of tactical aircraft sequencing in a model with normally-distributed aircraft approach speeds and runway occupancy times. Gupta et al. (2011) (see also Atkin et al. (2008), Xue and Zelinski (2015), Matsuno et al. (2017)) have tested the performances of deterministic schedule optimization methods under demand-related uncertainty. It should be noted that the question of how to make dynamic updates to predicted aircraft arrival times in response to the latest available information is an important area of study in its own right (Bronsvoort et al. (2009), Tobaruela et al. (2014), Tielrooij et al. (2015)). Niendorf et al. (2016) have discussed the use of stability analysis in order to quickly detect whether or not an optimized aircraft landing sequence remains optimal after unforeseen delays have occurred. Notably, their approach can also be used to determine whether or not a first-come-first-served sequence is optimal.

In the last decade, some promising stochastic optimization approaches have been developed for RSPs. Solveling et al. (2011) developed a two-stage approach in which a sequence of aircraft weight classes is determined in the first stage, and individual flights are assigned to positions in the sequence (subject to weight class compatibility) after operational uncertainty has been realized. Bosson et al. (2015) proposed a formulation that allows for greater integration with the problem of scheduling surface operations, with additional uncertainties related to push-back and taxiing times. Solak et al. (2018) extended the earlier work of Solveling et al. (2011) by enhancing the second stage problem so that total costs are based on exact timings of runway movements (as opposed to aircraft sequence positions). The uncertainty in their model is related to the times at which flights become eligible to be added to the runway sequence. Two-stage stochastic optimization methods have also been employed in computational studies by Liu et al. (2018) and Khassiba et al. (2020). In realistic-sized problems, it is invariably necessary to employ the sample average approximation (see Kleywegt et al. (2002)), in order to restrict the set of scenarios under consideration.

Although stochastic optimization methods have received considerable recent attention, other approaches based on robust optimization are possible. Heidt et al. (2016) compared ‘strict robust’ and ‘light robust’ approaches in a problem where dynamic random perturbations are applied to the lower and upper bounds of the time interval in which a particular aircraft is able to take off or land. Strict robustness requires an operating plan to be feasible under all possible random scenarios, but light robustness allows the user to sacrifice a certain amount of stability (i.e. protection against enforced schedule changes) in order to achieve a better performance with respect to throughput rates and delays. Heidt et al. (2016) used simulation experiments to show that these methods are effective in achieving the desired trade-off between stability and efficiency. Ng et al. (2017) also used robust optimization methods in a problem that included runway configuration planning; specifically, they employed a min-max regret approach.

As stochastic and robust optimization methods have gained traction in the runway scheduling literature, insights into the VSS have become possible. As in other problem contexts (see our discussions in Sections 3.1

and 4.1), the sensitivity of the VSS to parameter changes and other model adjustments is often worthwhile to investigate. Solveling et al. (2011) and Solak et al. (2018) found that the benefits of stochastic optimizers tend to increase as flight schedules become more densely populated. Khassiba et al. (2020) found that the VSS could be greater than 10% in some of their problem instances, but the inclusion of chance constraints related to aircraft separation standards can have a (possibly) surprising effect. Specifically, if α represents the minimum acceptable probability of separation standards being maintained by two consecutive aircraft in a landing sequence, then the VSS falls sharply as α increases towards one. This is because high values of α cause elongated separations to be required in order to obtain the necessary protection levels against uncertainty. In these circumstances, landing sequences are less likely to be disrupted as events unfold, and the benefits of using a stochastic optimizer are therefore reduced.

In the runway scheduling context, the benefits of robust optimization methods are witnessed in the form of fewer enforced sequence changes. Heidt et al. (2016) found that the number of ‘go-arounds’ for arrivals or slot losses for departures (resulting from disturbances to planned runway sequences) could be reduced to zero, even under the ‘light robustness’ version of their model. Usually, the use of robust optimization should cause the duration or ‘makespan’ of a runway sequence to increase due to the longer time separations required, but the opposite effect can occur if non-robust models are forced to move individual flights to later positions in the sequence due to their planned operation times becoming infeasible. A key assumption made by Heidt et al. (2016) (which also applies to the stochastic optimization methods discussed earlier) is that the decision-maker has explicit knowledge of the relevant uncertainty distributions. Clearly, the effectiveness of a robust or stochastic optimization approach is likely to be severely compromised if the nature of the uncertainty is unknown. We discuss this further in our conclusions.

It is interesting to note that the literature on RSPs has developed along quite separate lines from that of the airport capacity utilization problems discussed in Sections 3.2 and 3.3. Notably, while the aforementioned capacity utilization problems are increasingly making use of stochastic queueing models to represent operational uncertainty, the treatment of uncertainty in RSPs tends to be more limited; indeed, it is only in recent years that non-deterministic solution methods for these problems have been properly explored. Undoubtedly this is due to the more complicated modeling assumptions (with respect to individual flight attributes, multiple objectives, etc.) inherent in RSPs, which lend themselves more to stochastic integer programming formulations. Nevertheless, it appears that there is potential for greater synthesis between these different types of problems. Indeed, the fundamental objective in RSPs is to optimize the composition of an aircraft queue (or multiple queues), so this suggests that there should be interesting possibilities for incorporating stochastic queueing dynamics to a greater extent in such problems. It is clear, however, that the computational challenges involved are not to be underestimated.

4.3 Aircraft conflict detection and resolution

The foremost responsibility of pilots and air traffic controllers is to ensure that aircraft arrive at their destinations safely. The problem of identifying potentially dangerous situations caused by conflicts between the flight trajectories of two or more aircraft, and prescribing avoidance measures in such a way as to (ideally) minimize the amount of disruption to aircraft itineraries and routes, is referred to in the

literature as the aircraft conflict detection and resolution (CDR) problem.

Although one might tend to imagine CDR problems as being operational problems defined on short time horizons, in fact the CDR literature is quite diverse and the relevant issues have also been incorporated into strategic planning processes (Sherali et al. (2003, 2006), Netjasov (2012), Courchelle et al. (2019)). In a recent literature survey, Tang (2019) classified CDR problems into three different categories, depending on the decision-making horizon of interest. ‘Long-term’ CDR takes place in advance of flight execution, and involves the strategic coordination of aircraft trajectories in order to reduce potential conflicts whilst also respecting air sector capacities and other airspace restrictions. ‘Medium-term’ CDR incorporates real-time information about the positions, speeds etc. of aircraft in an airspace region and offers tactical solutions (including heading, speed and altitude changes) that enable dangerous situations to be prevented. ‘Short-term’ CDR, as the name suggests, addresses critical situations in which immediate maneuvers are required by aircraft in order to ensure safety. All three types of problems have been approached using various different mathematical prediction and optimization techniques, and many (albeit not all) of the relevant formulations aim to model the effects of uncertainties such as meteorological conditions, instrument precision, speed of human response to instructions, etc. In keeping with the scope of the paper, this subsection focuses on probabilistic methods for CDR, although many important contributions have been made by authors who considered deterministic settings (Bicchi and Pallottino (2000), Pallottino et al. (2002), Alonso-Ayuso et al. (2011, 2016), Cafieri and Durand (2014), Omer (2015)).

Probabilistic approaches to CDR problems received very little attention until the 1990s (Kuchar and Yang (2000)). Prior to this, ‘worst-case’ approaches were used by some authors as a means of handling uncertainty (Ratcliffe (1989), Ford and Powell (1990)). Simply put, worst-case approaches consider a range of possible maneuvers for all aircraft in a particular region, and predict a conflict if any combination of these maneuvers (however unlikely) results in the pre-determined separation standards being violated. Although the motivation for such a conservative approach is obvious, in medium-term problems it is neither desirable nor realistic to reduce the probability of a conflict to zero (Erzberger et al. (1997)). Indeed, worst-case approaches can severely overestimate the number of conflicts that will occur, which results in unmanageable workloads for air traffic controllers (Alliot et al. (2001), Archambault (2004), Rey et al. (2016)). Probabilistic approaches are somewhat more flexible and allow for a more finely-judged assessment of the timing and extent of interventions that should be required in order to reduce conflict probabilities to reasonably low levels, taking into account the abilities of air traffic controllers and flight management systems to resolve conflicts if they do occur.

The essence of a probabilistic method is to estimate the probability of separation standards being violated. This can inform the specification of an ‘alert zone’, represented as a virtual cylinder around an aircraft at a given point in time, such that encroachment into the alert zone by another aircraft would trigger a ‘potential conflict’ warning (Tomlin et al. (1998), Kuchar and Yang (2000)). Estimation of conflict probabilities poses a formidable modeling challenge due to the various sources of uncertainty that must be considered. Erzberger et al. (1997) used an ellipsoidal region in three-dimensional space to represent the prediction error associated with an aircraft’s future trajectory, with the longest axis of the ellipsoid being in the ‘along-track’ direction. They also modeled these prediction errors as being normally distributed, which

results in the ellipsoidal regions being ‘stretched’ in the along-track direction as the prediction horizon becomes longer; see Figure 6. A geometrical method can then be used to calculate conflict probabilities based on the interactions between ellipsoidal regions corresponding to pairs of aircraft; see (Erzberger et al. (1997), Paielli and Erzberger (1997), Paielli (1998)) for details.




Figure 6: Trajectory prediction error ellipsoids, adapted from Erzberger et al. (1997)

Various other authors have also used Gaussian distributions to model the uncertainties in predicted aircraft trajectories due to speed variations, wind effects, etc. (Krozel and Peters (1997), Prandini et al. (2000), Irvine (2002), Lygeros and Prandini (2002), Hu et al. (2005), Matsuno et al. (2015)). In addition, Ballin and Erzberger (1996) and Wanke (1997) have provided empirical evidence to support the assumption of normally distributed along-track errors. An advantage of using Gaussian distributions is that one can easily model the accumulation of errors from multiple sources, so they result in more tractable methods of analysis. For example, one possible approach for evaluating the risk of a conflict between two aircraft during a finite horizon of length T (see, for example, Prandini et al. (2000)) is to calculate

$$\max_{t \in [0, T]} \int_{y \in C} p_{d_t}(y) dy,$$

where p_{d_t} is the probability density function for the separation distance between the aircraft at time t , and C is the set of values that would indicate a conflict. If the positions of the two aircraft are uncorrelated normal random variables, then p_{d_t} is simply a Gaussian density. As noted in Prandini et al. (2000), however, the assumption of uncorrelated position vectors is somewhat unrealistic in practice, as tracking errors are often influenced by wind conditions.

Extensions of Gaussian models are also possible. Jilkov et al. (2014), building upon the work of Blom and Bakker (2002), argued that a Gaussian mixture (GM) distribution is appropriate to use in a multiple model trajectory prediction framework such as NextGen, which allows for consideration of many different aircraft maneuvers (including turning, climbing and descent, etc.). In a GM distribution, the probability density function is obtained as a weighted sum of Gaussian densities and (in general) is not Gaussian itself. Jilkov et al. (2014) proposed an efficient numerical method for evaluating the conflict probability in a model with GM-distributed separation vectors between aircraft. In addition, the modeling of wind conditions is often underpinned by Gaussian processes, although it is important to take correlation structure into account in such models. A typical approach, as used in Chaloulos et al. (2010), is to model wind velocity as the sum of a deterministic (nominal) component and a stochastic component, where the stochastic component has the form of a Gaussian random field. This type of model can ensure wind correlation in time and space; in

other words, the wind conditions experienced by an aircraft at a particular time are correlated with those of other aircraft at the same time, and also with the wind conditions at earlier times. If such correlation effects are ignored, then the result is a ‘white noise’ process, which may be unreliable for estimating conflict probabilities (Chaloulos and Lygeros (2007)).

So far, the discussion in this subsection has focused on modeling issues and the estimation of conflict probabilities. We now proceed to discuss decision-making and optimization problems. Several authors have considered the problem of directing aircraft to their destinations in an expeditious manner while maintaining acceptable bounds on the probability of conflict, and this has been treated as a stochastic optimal control problem; see Lecchini Visintini et al. (2006), Kantas et al. (2010), Liu and Hwang (2014), Matsuno et al. (2015), Hentzen et al. (2018). In Kantas et al. (2010), for example, the objective is to minimize the expected maximum time of arrival among a group of aircraft while ensuring a lower bound on the probability of separation standards being maintained. From a computational perspective, such problems involve expectations over high-dimensional probability distributions, and Monte Carlo approximations are often required. Liu and Hwang (2014) incorporated a stochastic differential equation in order to model wind and weather variations, and approximated its behavior using a Markov chain defined on a discretized version of the state space. The resulting transition equations are solved using the Jacobi method. Their method, although impressive, is also somewhat computationally demanding (Jilkov et al. (2018)). In Hentzen et al. (2018), the problem of interest is to guide aircraft to waypoints while avoiding hazardous storm regions. This problem belongs to the class of *stochastic reachability* (or *reach-avoid*) decision problems, which have been studied in more general settings (Summers and Lygeros (2010), Esfahani et al. (2016)) and have found other applications in CDR problems (Watkins and Lygeros (2003), Yang et al. (2017)).

Applications of other optimization techniques are somewhat more difficult to find in the CDR literature, although examples do exist. Vela et al. (2009) used a two-stage stochastic program with recourse to address the problem of assigning speed changes to aircraft in order to avoid conflicts. The first stage takes place before wind uncertainty (modeled using normally-distributed random variables) is realized, and the second stage prescribes last-minute maneuvers if actual wind conditions imply that safety is not guaranteed. The objective is based on minimization of fuel costs. Although the results appear encouraging, it may be necessary to develop more sophisticated models in order to handle complex trajectory uncertainties (Allignol et al. (2013), Wang et al. (2020)). Rey et al. (2016) formulated deterministic optimization models with objectives related to ATC workload (e.g. total number of conflicts), but then used a method similar to that of Haddad et al. (2008) to ‘stress’ their solutions in a stochastic environment with aircraft speed perturbations. Lehouillier et al. (2017) approached the CDR problem from a multi-objective perspective and developed a decision analysis tool that presents controllers with a set of possible solutions, based on the trade-offs between fuel consumption, ATC workload and other relevant criteria. Their iterative optimization algorithm involves the solution of maximum clique problems (formulated as MILPs) and includes consideration of various uncertainty sources which are modeled using Gaussian distributions.

In summary, CDR problems (like many of the other types of problems considered in this survey) pose their own unique set of modeling challenges, particularly where uncertainty is concerned. Many of the older papers in this area focus on prescribing maneuvers to pairs of aircraft in order to avoid conflicts, but as

the ATM landscape evolves (with the progression of projects such as NextGen and SESAR) it will become increasingly important to consider CDR issues at a more strategic level and consider the implications for air sector capacities, feasible ATC workloads, etc. There are obvious links between CDR problems and the ATFM problems discussed in Section 4.1, but to date it appears that relatively little exploration has been done into integrating the relevant objectives and solution methods.

5 Discussion and conclusions

The previous sections have described a wide range of stochastic modeling applications in problems related to demand and capacity management at airports and airport networks and identified new and exciting possibilities for future research. This section summarizes the key findings of this survey and discusses how new research ideas might be taken forward. The following questions are considered:

- How is the air transport research landscape likely to change in the future?
- What impact will these changes have on the mathematical formulations and solution approaches commonly employed in the current literature?
- What are the most promising future research opportunities for stochastic modelers?

These questions are addressed in the subsections that follow.

5.1 Changes to the research landscape in the ‘era of big data’

Firstly, it may be observed that the continuous evolution of computing power and the increasing availability of real-world data have enabled powerful new research methods that would not have been considered practical or feasible in previous decades. Recently, Li and Ryerson (2019) presented a review of articles published since 2010 in order to show how air transport research is changing in the era of ‘big data’. Indeed, it should be noted that the use of data mining, forecasting and machine learning algorithms (to predict flight delays, for example) is gaining popularity in the air transport research community (Deshpande and Arikan (2012), Barnhart et al. (2014), Rebollo and Balakrishnan (2014), Hanley (2015), Choi et al. (2016), Gopalakrishnan and Balakrishnan (2017), Belcastro et al. (2018), Munoz et al. (2018)).

This paper has not attempted to cover machine learning methods in detail because they do not easily fit within the scope of our literature survey. However, the interface between artificial intelligence (AI) and optimization is becoming more critical and AI methods have great potential for cross-validating the results of optimization procedures or being integrated within hybridized decision support systems. It is now common for AI methods to be used for *prescriptive* (as well as predictive) purposes; for example, Estes et al. (2018) described the use of a novel unsupervised learning method for informing air traffic management decisions in the context of ground delay programs.

Although unsupervised learning and other AI algorithms are clearly versatile enough to find meaningful patterns in datasets that have been affected by multiple sources of uncertainty, they also have strong potential for validating the assumptions inherent in more traditional OR model designs (e.g. queuing

system formulations for air traffic). From a stochastic modeling perspective, there are exciting possibilities for AI methods to assist model designs by ensuring that uncertainty is modeled in a way that accords with real-world experience. As noted in Section 4.2, robust and stochastic optimization methods tend to rely on explicit and accurate knowledge of the relevant uncertainty distributions being available, so this creates opportunities for the integration of machine learning algorithms.

Data-driven methods are already making an impact in queueing system models for air traffic. It is clear that computationally tractable models such as $M(t)/E_k(t)/1$ continue to enjoy strong popularity; indeed, the $M(t)/E_k(t)/1$ model is quite versatile due to its time-dependent demand and service rates and the ability of Erlang distributions to closely approximate many parametric and empirical distributions. Furthermore, the numerical approximation methods discussed in Section 2.4 (e.g. DELAYS, SBC) enable the efficient computation of time-dependent probability distributions without time-consuming Monte Carlo simulations being required. However, these ‘classical’ queueing formulations are likely to come under increasing scrutiny as the availability of real-time traffic data makes it easier to challenge their assumptions. Data-driven pre-scheduled random demand (PSRD) models are becoming more popular in the literature, and existing computational studies have shown that these compare favorably to Poisson models with respect to prediction of queue lengths and delays (Caccavale et al. (2014), Gwiggner and Nagaoka (2014), Lancia and Lulli (2020)).

We suggest that, as one possible direction for future research, it may be possible to investigate the theoretical properties of certain classes of PSRD models and develop numerical approximation algorithms (similar in purpose to those that exist for Poisson models) in order to compute more reliable estimates for time-dependent performance measures in air traffic queues.

5.2 Stochastic modeling as a tool to inform strategic decision-making

Slot allocation problems, discussed in Section 3.1, feature the control of airport demand rates as a prominent theme. The trade-off between schedule displacement and operational delays has been well-observed in the literature (Barnhart et al. (2012), Swaroop et al. (2012)), and this has highlighted the importance of setting slot controls (interpreted as capacity constraints in optimization problems) appropriately at slot-coordinated airports in order to place restrictions on expected flight delays.

An important principle is that if flight delays are predicted using *deterministic* queueing dynamics (which can easily be incorporated within integer programming formulations for slot allocation), then the resulting predictions are likely to be overly optimistic, and any slot allocation mechanism based on these predictions is likely to allow too many flights to be scheduled within short time intervals. It is only by modeling flight delays *stochastically* that one can gain accurate forecasts of expected congestion levels and delays. This principle has been employed to great effect by Jacquillat and Odoni (2015a), whose seminal paper proposed a slot allocation framework that iterates between an integer programming model for slot allocation and a stochastic dynamic programming model (with $M(t)/E_k(t)/1$ queue dynamics) for capacity utilization in order to optimize the trade-off between scheduling and operational delays.

In a sense, the model of Jacquillat and Odoni (2015a) circumvents the need for slot controls because it evaluates the operational feasibility (i.e. queueing performance) of a particular schedule using a dynamic,

stochastic model of capacity utilization, rather than ensuring that the schedule satisfies an exogenous set of slot capacity constraints. Given that, in practice, airport slot controls are usually determined at an administrative level rather than being informed by any form of stochastic modeling (see Zografos et al. (2017)), it is possible to argue that the approach of Jacquillat and Odoni (2015a) is more powerful and has the potential to offer ‘better’ schedules (with respect to the trade-off between conflicting objectives) than those that would be obtained by imposing a set of slot constraints *ex ante*. In reality, however, the picture is more complicated. The strategy of imposing administrative slot controls (e.g. a maximum of 40 flights per hour) is well-understood by industry practitioners, and to abandon this system in favor of a more sophisticated approach based on stochastic queue modeling would require a significant amount of trust to be placed in the validity of the underlying modeling assumptions; for example, the queueing model in Jacquillat and Odoni (2015a) is based on $M(t)/E_k(t)/1$ dynamics which (as discussed elsewhere in this paper) have been questioned by several researchers. We suggest, therefore, that the problem of trying to determine an optimal set of slot capacity constraints (considered by Churchill et al. (2013)) by modeling flight delays stochastically remains worthy of attention.

There are also possibilities for considering similar problems at a network level. Some progress has been made in formulating network-level slot allocation problems in the last few years (Pellegrini et al. (2017)), and the possible synergies between these and ATFM problems are beginning to be explored. Wang and Jacquillat (2020) considered a US-centric problem in which strategic scheduling interventions are followed by tactical ground-holding decisions, with the latter (but not the former) taking place after the realization of weather-related uncertainty. However, their stochastic programming formulation does not include air sector capacities or fine-grain models of queueing dynamics at individual airports. Furthermore, in order to consider a similar problem in a European context, one would need to include constraints based on a scheduling season rather than a single day of operations. We anticipate that large-scale stochastic programming formulations will be developed in the coming years to address some of these open problems.

Simulation-based optimization methods may offer a way forward in problems where the interactions between different sources of uncertainty are difficult to capture using analytically tractable models. Techniques based on adaptive random search or gradient descent (see Nelson (2013), Fu (2015)), which would have been considered computationally infeasible in the past, have proven themselves capable of finding strong-performing solutions in problems with vast solution spaces in which solution ‘strength’ can only be estimated using Monte Carlo methods. In a network-level slot allocation problem, for example, one might use high-fidelity simulation experiments to estimate the expected ATFM-related delays under various ‘candidate’ schedules and eliminate those that do not achieve the required robustness standards.

The type of approach discussed above should be feasible because slot allocations are produced at the *strategic* level, so one can allow plenty of CPU running time to obtain high-quality solutions. In tactical problems where decisions must be made dynamically in real time, it is less clear whether or not such methods are practical.

5.3 Fast solution approaches for rapidly-changing problem environments

Multi-stage stochastic optimization methods have the potential to make a strong impact in problems related to airport resource allocation and traffic flow management. The relevant techniques in this area include stochastic dynamic programming (Ross (1983), Puterman (2005)), approximate dynamic programming (Bertsekas and Tsitsiklis (1996), Powell (2007)) and stochastic nested decomposition (Birge and Louveaux (2011), Kall and Mayer (2011)). Since these problems require decisions to be made in response to the latest unfolding events on a day of operations, it is natural to adopt a Markov decision process (MDP) formulation. Some progress has been made in solving airport capacity utilization problems with low-dimensional state and action spaces (Jacquillat et al. (2017), Shone et al. (2019)), but there may be possibilities for applying ADP methods to higher-dimensional problems.

In problems related to ATFM or aircraft conflict detection and resolution (CDR), the information available to a decision-maker should include detailed information about the latest positions and estimated waypoint arrival times of individual flights as well as forecasts of future weather and wind conditions, and decisions should be made with some appreciation or (preferably) explicit modeling of possible random variations over the time horizon of interest. In CDR problems, methods from stochastic optimal control theory have already been explored (Liu and Hwang (2014), Matsuno et al. (2015)), but these tend to rely upon simplified representations of state and action spaces. ADP methods that employ value function approximation and feature extraction may offer practical solutions to problems with vast spaces, and this may facilitate the formulation of optimization problems which combine commonly-used objectives from the ATFM and CDR literatures. We also suggest that decomposition algorithms from the wider stochastic optimization literature (see Escudero et al. (2012, 2016), Zou et al. (2019)) are worthy of attention in such problems.

Runway scheduling problems (RSPs) are clearly related to the dynamic optimization problems described above, but until now these have largely been treated as a separate class of problem in the literature. This is because the assignment of specific attributes (e.g. weight class, preferred landing time, etc.) to individual flights naturally places such problems within the realm of combinatorial optimization. Considerable progress has been made in recent years on developing stochastic and robust optimization models for RSPs (Heidt et al. (2016), Solak et al. (2018), Khassiba et al. (2020)), but the types of uncertainty included in such formulations to date have been somewhat limited. More specifically, uncertainty is usually introduced with respect to the earliest and latest permissible take-off/landing times of aircraft, but not with respect to aircraft 'service' (i.e. inter-landing or inter-departure) times.

The use of Erlang distributions (in particular) to model aircraft service times has been widely adopted in other areas of the literature, and these distributions can potentially be calibrated according to different 'leader-follower' pairs of aircraft types, as described by Jeddi et al. (2006). We therefore suggest that new types of RSPs that incorporate multiple sources of uncertainty may be interesting to study.

5.4 Emerging opportunities for stochastic modeling

To summarize the discussion in this section, we suggest that some of the most exciting future research opportunities for stochastic modelers in air traffic management include the following:

- The use of machine learning and other AI methods to fine-tune the parameters and probability distributions used in stochastic models according to empirical data, or to cross-validate the solutions found using stochastic optimization methods;
- The development of new, computationally efficient numerical methods for approximating key, time-dependent performance measures in queueing systems with complicated dynamics (e.g. those with non-Markovian transitions), including network extensions;
- The use of stochastic programming to develop new problem formulations that directly incorporate multiple sources/types of uncertainty, e.g. by including constraints or objectives based on queueing performance measures with nonlinear behavior;
- The development of innovative solution algorithms for reducing the complexity of very large stochastic optimization formulations, e.g. via decomposition methods;
- The use of simulation-based optimization methods to evaluate the performances of candidate solutions in highly stochastic environments and identify those that are sufficiently robust;
- The design and implementation of ADP methods for obtaining strong-performing policies in dynamic problems with high-dimensional state and action spaces;
- The development of adaptable optimization models that can easily incorporate the objectives of multiple stakeholders, including (for example) those which require risk averse modeling.

In conclusion, applications of stochastic modeling in air traffic management are continuously evolving. The research agenda is being shaped by many different factors, including (i) changes in strategic and operational practices within the air transport system, (ii) the increasing diversity and availability of aviation-related data, (iii) the growing potential for computing procedures to find optimal or strong-performing solutions in large-scale decision-making problems. This survey has focused on a small number of specific stochastic modeling techniques that offer a wide range of potential applications. We anticipate that similar techniques will continue to find further applications in the coming years, but we also look forward to the emergence of new and innovative solution methods that may be required to tackle the next generation of research problems in air traffic management.

Acknowledgements

We would like to thank Professor Amedeo Odoni for his generous advice and support during the preparation of this paper. We would also like to thank the anonymous referees for their valuable comments and suggestions.

This work has been supported by the Engineering and Physical Sciences Research Council (EPSRC) through Programme Grant EP/M020258/1 “Mathematical models and algorithms for allocating scarce airport resources (OR-MASTER)”.

References

- Abundo, S. (1990). *An Approach for Estimating Delays at a Busy Airport*. Master's Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Agustin, A., Alonso-Ayuso, A., Escudero, L., and Pizarro, C. (2010). *Mathematical optimization models for Air Traffic Flow Management: A review*. In: Bui, A. and Tseveendorkj, I. (eds). *Combinatorial Optimization in Practice*, Studia Informatica Universalis, Hermann Informatique, 8(2).
- Agustin, A., Alonso-Ayuso, A., Escudero, L., and Pizarro, C. (2012a). On air traffic flow management with rerouting. Part I: Deterministic case. *European Journal of Operational Research*, 219:156–166.
- Agustin, A., Alonso-Ayuso, A., Escudero, L., and Pizarro, C. (2012b). On air traffic flow management with rerouting. Part II: Stochastic case. *European Journal of Operational Research*, 219:167–177.
- Allignol, C., Barnier, N., Durand, N., and Alliot, J.-M. (2013). A New Framework for Solving En Route Conflicts. *Air Traffic Control Quarterly*, 21(3):233–253.
- Alliot, J.-M., Durand, N., and Granger, G. (2001). *A statistical analysis of the influence of vertical and ground speed errors on conflict probe*. Proc. 4th USA/Europe Air Traffic Management R&D Seminar, Santa Fe, NM.
- Alonso, A., Escudero, L., and Teresa Ortuno, M. (2000). A stochastic 0-1 program based approach for the air traffic flow management problem. *European Journal of Operational Research*, 120:47–62.
- Alonso-Ayuso, A., Escudero, L., and Martin-Campo, F. (2011). Collision avoidance in air traffic management: a mixed-integer linear optimization approach. *IEEE Transactions on Intelligent Transportation Systems*, 12(1):47–57.
- Alonso-Ayuso, A., Escudero, L., and Martin-Campo, F. (2016). Multiobjective optimization for aircraft conflict resolution. A metaheuristic approach. *European Journal of Operational Research*, 248:691–702.
- Andersson, K., Carr, F., Feron, E., and Hall, W. (2000). *Analysis and Modeling of Ground Operations at Hub Airports*. Proc. 3rd USA/Europe Air Traffic Management R&D Seminar, Napoli, Italy.
- Andreatta, G., Dell'Olmo, P., and Lulli, G. (2011). An aggregate stochastic programming model for air traffic flow management. *European Journal of Operational Research*, 215:697–704.
- Andreatta, G. and Lulli, G. (2009). Equitable Demand Management Strategies for Different Classes of Customers. *International Journal of Pure and Applied Mathematics*, 57(1):1–22.
- Andreatta, G. and Odoni, A. (2003). *Analysis of market-based demand management strategies for airports and en route airspace*. In: Ciriani, T.A., Fasano, G., Gliozzi, S. and Tadei, R. (eds). *Operations Research in Space and Air*. Kluwer.
- Andreatta, G. and Romanin-Jacur, G. (1987). Aircraft Flow Management under Congestion. *Transportation Science*, 21(4):249–253.
- Aoun, O. and El Afia, A. (2014). *Using Markov decision processes to solve stochastic gate assignment problem*. Proc. International Conference on Logistics Operations Management, Rabat, Morocco.
- Archambault, N. (2004). *Speed uncertainty and speed regulation in conflict detection and resolution in air traffic control*. Proc. 1st International Conference on Research in Air Transportation, Zilina, Slovakia.
- Arikan, M., Deshpande, V., and Sohoni, M. (2013). Building Reliable Air-Travel Infrastructure Using Empirical Data and Stochastic Models of Airline Networks. *Operations Research*, 61(1):45–64.

- Atkin, J., Burke, E., Greenwood, J., and Reeson, D. (2008). On-line decision support for take-off runway scheduling with uncertain taxi times at London Heathrow airport. *Journal of Scheduling*, 11:323–346.
- Badrinath, S. and Balakrishnan, H. (2017). *Control of a non-stationary tandem queue model of the airport surface*. Proc. 2017 American Control Conference (ACC), Seattle, WA.
- Badrinath, S., Balakrishnan, H., Joback, E., and Reynolds, T. (2020). Impact of Off-Block Time Uncertainty on the Control of Airport Surface Operations. *Transportation Science*, 54(4):920–943.
- Balakrishna, P., Ganesan, R., and Sherry, L. (2010). Accuracy of reinforcement learning algorithms for predicting aircraft taxi-out times: A case-study of Tampa Bay departures. *Transportation Research Part C*, 18:950–962.
- Balakrishnan, H. and Chandran, B. (2010). Algorithms for scheduling runway operations under constrained position shifting. *Operations Research*, 58(6):1650–1665.
- Balakrishnan, H. and Chandran, B. (2014). Optimal large-scale air traffic flow management. http://web.mit.edu/Hamsa/www/pubs/BalakrishnanChandran_ATFM.pdf. Accessed on Feb 22, 2019.
- Ball, M., Barnhart, C., Dresner, M., Hansen, M., Neels, K., Odoni, A., Peterson, E., Sherry, L., Trani, A., and Zou, B. (2010a). *Total delay impact study: A comprehensive assessment of the costs and impacts of flight delay in the United States*. Technical report, Federal Aviation Administration, Washington, DC.
- Ball, M., Donohue, G., and Hoffman, K. (2006). *Auctions for the safe, efficient and equitable allocation of airspace system resources*. In: Cramton, P., Shoham, Y. and Steinberg, R. (eds). *Combinatorial Auctions*. Cambridge: MIT Press.
- Ball, M., Estes, A., Hansen, M., and Liu, Y. (2020). Quantity-Contingent Auctions and Allocation of Airport Slots. *Transportation Science*, 54(4):858–881.
- Ball, M., Hoffman, R., Knorr, D., Wetherly, J., and Wambsganss, M. (2000). *Assessing the Benefits of Collaborative Decision Making in Air Traffic Management*. Proc. 3rd USA/Europe Air Traffic Management R&D Seminar, Napoli, Italy.
- Ball, M., Hoffman, R., and Mukherjee, A. (2010b). Ground Delay Program Planning Under Uncertainty Based on the Ration-by-Distance Principle. *Transportation Science*, 44(1):1–14.
- Ball, M., Hoffman, R., Odoni, A., and Rifkin, R. (2003). A Stochastic Integer Program with Dual Network Structure and its Application to the Ground-Holding Problem. *Operations Research*, 51(1):167–171.
- Ball, M., Vossen, T., and Hoffman, R. (2001). *Analysis of Demand Uncertainty Effects in Ground Delay Programs*. Proc. 4th USA/Europe Air Traffic Management R&D Seminar, Santa Fe, NM.
- Ballin, M. and Erzberger, H. (1996). *An analysis of landing rates and separations at the Dallas/FortWorth International Airport*. Technical report, Ames Research Center, Moffett Field, CA.
- Barnhart, C., Belobaba, P., and Odoni, A. (2003). Applications of Operations Research in the Air Transport Industry. *Transportation Science*, 37(4):368–391.
- Barnhart, C., Fearing, D., Odoni, A., and Vaze, V. (2012). Demand and capacity management in air transportation. *EURO Journal on Transportation and Logistics*, 1:135–155.
- Barnhart, C., Fearing, D., and Vaze, V. (2014). Modeling Passenger Travel and Delays in the National Air Transportation System. *Operations Research*, 62(3):580–601.
- Baspinar, B., Kemal Ure, N., Koyuncu, E., and Inalhan, G. (2016). Analysis of Delay Characteristics of European Air Traffic through a Data-Driven Airport-Centric Queuing Network Model. *IFAC-PapersOnLine*, 49(3):359–364.

- Basso, L. and Zhang, A. (2010). Pricing vs. slot policies when airport profits matter. *Transportation Research Part B*, 44(3):381–391.
- Bauerle, N., Engelhardt-Funke, O., and Kolonko, M. (2007). On the waiting time of arriving aircrafts and the capacity of airports with one or two runways. *European Journal of Operational Research*, 177(2):1180–1196.
- Beasley, J., Krishnamoorthy, M., Sharaiha, Y., and Abramson, D. (2004). Displacement Problem and Dynamically Scheduling Aircraft Landings. *Journal of the Operational Research Society*, 55(1):54–64.
- Belcastro, L., Marozzo, F., Talia, D., and Trunfio, P. (2018). Using Scalable Data Mining for Predicting Flight Delays. *ACM Transactions on Intelligent Systems and Technology*, 8(1).
- Bennell, J., Mesgarpour, M., and Potts, C. (2017). Dynamic scheduling of aircraft landings. *European Journal of Operational Research*, 258:315–327.
- Bertsekas, D. and Tsitsiklis, J. (1996). *Neuro-Dynamic Programming*. Athena Scientific.
- Bertsimas, D., Frankovich, M., and Odoni, A. (2011a). Optimal selection of airport runway configurations. *Operations Research*, 59(6):1407–1419.
- Bertsimas, D., Lulli, G., and Odoni, A. (2011b). An Integer Optimization Approach to Large-Scale Air Traffic Flow Management. *Operations Research*, 59(1):211–227.
- Bertsimas, D. and Stock Patterson, S. (1998). The Air Traffic Flow Management Problem with Enroute Capacities. *Operations Research*, 46(3):406–422.
- Bicchi, A. and Pallottino, L. (2000). On Optimal Cooperative Conflict Resolution for Air Traffic Management Systems. *IEEE Transactions on Intelligent Transportation Systems*, 1(4):221–232.
- Birge, J. and Louveaux, F. (2011). *Introduction to Stochastic Programming, 2nd edition*. Springer.
- Blom, H. and Bakker, G. (2002). *Conflict probability and incrossing probability in air traffic management*. Proc. 41st IEEE Conference on Decision and Control, Las Vegas, NV.
- Blumstein, A. (1959). The Landing Capacity of a Runway. *Operations Research*, 7(6):752–763.
- Bolender, M. and Slater, G. (2000). Evaluation of scheduling methods for multiple runways. *Journal of Aircraft*, 37(3):410–416.
- Bookbinder, J. (1986). Multiple Queues Of Aircraft Under Time-Dependent Conditions. *INFOR: Information Systems and Operational Research*, 24:280–288.
- Bosson, C., Xue, M., and Zelinski, S. (2015). *Optimizing Integrated Arrival, Departure and Surface Operations Under Uncertainty*. Proc. 11th USA/Europe Air Traffic Management R&D Seminar, Lisbon, Portugal.
- Bouras, A., Ghaleb, M., Suryahatmaja, S., and Salem, A. (2014). The Airport Gate Assignment Problem: A Survey. *The Scientific World Journal*, pages 1–27.
- Bratu, S. and Barnhart, C. (2006). Flight operations recovery: New approaches considering passenger recovery. *Journal of Scheduling*, 9:279–298.
- Bronsvort, J., McDonald, G., Porteous, R., and Gutt, E. (2009). *Study of Aircraft Derived Temporal Prediction Accuracy using FANS*. Proc. 13th Air Transport Research Society World Conference, Abu Dhabi, UAE.
- Brueckner, J. (2009). Price vs. quantity-based approaches to airport congestion management. *Journal of Public Economics*, 93(5-6):681–690.

- Bureau of Transportation Statistics (2020). Understanding the Reporting of Causes of Flight Delays and Cancellations. <https://www.bts.gov/topics/airlines-and-airports/understanding-reporting-causes-flight-delays-and-cancellations>. Accessed on July 30, 2020.
- Burgain, P., Pinon, O., Feron, E., Clarke, J.-P., and Mavris, D. (2009). *On the value of information within a collaborative decision making framework for airport departure operations*. Proc. IEEE/AIAA 28th Digital Avionics Systems Conference, Orlando, FL.
- Buxi, G. and Hansen, M. (2011). *Generating Probabilistic Capacity Profiles from weather forecast: A design-of-experiment approach*. Proc. 9th USA/Europe Air Traffic Management R&D Seminar, Berlin, Germany.
- Caccavale, M., Iovanella, A., Lancia, C., Lulli, G., and Scoppola, B. (2014). A model of inbound air traffic: The application to Heathrow airport. *Journal of Air Transport Management*, 34:116–122.
- Cafieri, S. and Durand, N. (2014). Aircraft deconfliction with speed regulation: new models from mixed-integer optimization. *Journal of Global Optimization*, 58:613–629.
- Carr, F. (2004). *Robust decision-support tools for airport surface traffic*. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Chaloulos, G., Cruck, E., and Lygeros, J. (2010). A simulation based study of subliminal control for air traffic management. *Transportation Research Part C*, 18:963–974.
- Chaloulos, G. and Lygeros, J. (2007). Effect of wind correlation on aircraft conflict probability. *Journal of Guidance, Control and Dynamics*, 30(6):1742–1752.
- Chang, Y.-H., Solak, S., Clarke, J.-P., and Johnson, E. (2016). Models for single-sector stochastic air traffic flow management under reduced airspace capacity. *Journal of the Operational Research Society*, 67:54–67.
- Chen, H. and Solak, S. (2020). Lower cost departures for airlines: Optimal policies under departure metering. *Transportation Research Part C*, 111:531–546.
- Chen, J., Chen, L., and Sun, D. (2017). Air traffic flow management under uncertainty using chance-constrained optimization. *Transportation Research Part B*, 102:124–141.
- Choi, S., Kim, Y., Briceno, S., and Mavris, D. (2016). *Prediction of weather-induced airline delays based on machine learning algorithms*. Proc. IEEE/AIAA 35th Digital Avionics Systems Conference, Sacramento, CA.
- Churchill, A., Lovell, D., and Ball, M. (2010). Flight Delay Propagation Impact on Strategic Air Traffic Flow Management. *Transportation Research Record*, 2177:105–113.
- Churchill, A., Lovell, D., Mukherjee, A., and Ball, M. (2013). Determining the Number of Airport Arrival Slots. *Transportation Science*, 47(4):526–541.
- Churchill, A., Vlachou, K., and Lovell, D. (2008). *Filtering and aggregation schemes for delay model calibration*. Proc. 3rd International Conference on Research in Air Transportation (ICRAT), Fairfax, VA.
- Cieselski, V. and Scerri, P. (1997). An anytime algorithm for scheduling of aircraft landing times using genetic algorithms. *Australian Journal of Intelligent Information Processing Systems*, 4:206–213.
- Clare, G. and Richards, A. (2012). *Air traffic flow management under uncertainty: application of chance constraints*. Proc. 2nd International Conference on Application and Theory of Automation in Command and Control Systems, London, UK.
- Clark, C. (1961). The greatest of a finite set of random variables. *Operations Research*, 9(2):145–162.

- Clarke, J.-P., Melconian, T., Bly, E., and Rabbani, F. (2007). MEANS—MIT Extensible Air Network Simulation. *Simulation*, 83(5):385–399.
- Clarke, J.-P., Solak, S., Chang, Y.-H., Ren, L., and Vela, A. (2009). *Air Traffic Flow Management in the Presence of Uncertainty*. Proc. 8th USA/Europe Air Traffic Management R&D Seminar, Napa, CA.
- Clarke, J.-P., Solak, S., Ren, L., and Vela, A. (2013). Determining Stochastic Airspace Capacity for Air Traffic Flow Management. *Transportation Science*, 47(4):542–559.
- Corolli, L., Lulli, G., and Ntaimo, L. (2014). The time slot allocation problem under uncertain capacity. *Transportation Research Part C*, 46:16–29.
- Corolli, L., Lulli, G., Ntaimo, L., and Venkatachalam, S. (2017). A two-stage stochastic integer programming model for air traffic flow management. *IMA Journal of Management Mathematics*, 28:19–40.
- Courchelle, V., Soler, M., Gonzalez-Arribas, D., and Delahaye, D. (2019). A simulated annealing approach to 3D strategic aircraft deconfliction based on en-route speed changes under wind and temperature uncertainties. *Transportation Research Part C*, 103:194–210.
- Czerny, A. (2010). Airport congestion management under uncertainty. *Transportation Research Part B*, 44(3):371–380.
- Dal Sasso, V., Djeumou Fomeni, F., Lulli, G., and Zografos, K. (2018). Incorporating Stakeholders’ priorities and preferences in 4D trajectory optimization. *Transportation Research Part B*, 117:594–609.
- Dal Sasso, V., Djeumou Fomeni, F., Lulli, G., and Zografos, K. (2019). Planning efficient 4D trajectories in Air Traffic Flow Management. *European Journal of Operational Research*, 276:676–687.
- Daniel, J. (1995). Congestion Pricing and Capacity of Large Hub Airports: A Bottleneck Model with Stochastic Queues. *Econometrica*, 63(2):327–370.
- de Neufville, R. and Odoni, A. (2013). *Airport Systems: Planning, Design and Management, 2nd edition*. McGraw-Hill.
- Dear, R. (1976). *The Dynamic Scheduling of Aircraft in the Near Terminal Area*. Technical report R76-9, MIT Flight Transportation Laboratory.
- Defraeye, M. and Van Nieuwenhuyse, I. (2016). Staffing and scheduling under nonstationary demand for service: A literature review. *Omega*, 58:4–25.
- Dell’Olmo, P. and Lulli, G. (2003). A dynamic programming approach for the airport capacity allocation problem. *IMA Journal of Management Mathematics*, 14(3):235–249.
- Deshpande, V. and Arıkan, M. (2012). The Impact of Airline Flight Schedules on Flight Delays. *Manufacturing and Service Operations Management*, 14(3):423–440.
- Dorndorf, U., Jaehn, F., Lin, C., Ma, H., and Pesch, E. (2007). Disruption management in flight gate scheduling. *Statistica Neerlandica*, 61:92–114.
- Dunlay, W. and Horonjeff, R. (1976). Stochastic properties of enroute air traffic - An empirical investigation. *Journal of Aircraft*, 13(5):376–381.
- Erzberger, H., Paielli, R., Isaacson, D., and Eshow, M. (1997). *Conflict Detection and Resolution In the Presence of Prediction Error*. Proc. 1st USA/Europe Air Traffic Management R&D Seminar, Saclay, France.
- Escudero, L., Garin, M., Perez, G., and Unzueta, A. (2012). Lagrangian Decomposition for large-scale two-stage stochastic mixed 0-1 problems. *Top*, 20:347–374.

- Escudero, L., Garin, M., and Unzueta, A. (2016). Cluster Lagrangean decomposition in multistage stochastic optimization. *Computers & Operations Research*, 67:48–62.
- Esfahani, P., Chatterjee, D., and Lygeros, J. (2016). The stochastic reach-avoid problem and set characterization for diffusions. *Automatica*, 70:43–56.
- Estes, A. and Ball, M. (2020). Equity and Strength in Stochastic Integer Programming Models for the Dynamic Single Airport Ground-Holding Problem. *Transportation Science*, 54(4):944–955.
- Estes, A., Lovell, D., and Ball, M. (2018). Unsupervised prototype reduction for data exploration and an application to air traffic management initiatives. *EURO Journal on Transportation and Logistics*, pages 1–44.
- European Commission (2018). SESAR 2020: developing the next generation of European Air Traffic Management. https://ec.europa.eu/research/press/jti/factsheet_sesar-web.pdf. Accessed on May 21, 2018.
- Fairbrother, J. and Zografos, K. (2020). Optimal Scheduling of Slots with Season Segmentation. *European Journal of Operational Research (in press)*.
- Fairbrother, J., Zografos, K., and Glazebrook, K. (2020). A slot scheduling mechanism at congested airports which incorporates efficiency, fairness and airline preferences. *Transportation Science*, 54(1):115–138.
- Fan, T. (2003). *Market-based Airport Demand Management – Theory, Model and Applications*. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Fan, T. and Odoni, A. (2002). A practical perspective on airport demand management. *Air Traffic Control Quarterly*, 10:285–306.
- Ford, R. and Powell, D. (1990). A New Threat Detection Criterion for Airborne Collision Avoidance Systems. *The Journal of Navigation*, 43(3):391–403.
- Fu, M. (2015). *Handbook of Simulation Optimization*. Springer.
- Furstenau, N., Heidt, A., Kopolke, M., Liers, F., Mittendorf, M., and Weiss, C. (2015). *Pre-Tactical Planning of Runway Utilization Under Uncertainty: Optimization and Validation*. In: Schaefer, D. (ed), Proc. SESAR Innovation Days, Eurocontrol.
- Gallagher, H. and Wheeler, R. (1958). Nonstationary Queuing Probabilities for Landing Congestion of Aircraft. *Operations Research*, 6(2):264–275.
- Ganji, M., Lovell, D., Ball, M., and Nguyen, A. (2009). *Resource Allocation in Flow-Constrained Areas with Stochastic Termination Times Considering Both Optimistic and Pessimistic Reroutes*. Proc. 8th USA/Europe Air Traffic Management R&D Seminar, Napa, CA.
- Gilbo, E. (1993). Airport capacity: Representation, estimation, optimization. *IEEE Transactions on Control Systems Technology*, 1(3):144–154.
- Gilbo, E. (1997). Optimizing airport capacity utilization in air traffic flow management subject to constraints at arrival and departure fixes. *IEEE Transactions on Control Systems Technology*, 5(5):490–503.
- Gilbo, E. and Howard, K. (2000). *Collaborative optimization of airport arrival and departure traffic flow management strategies for CDM*. Proc. 3rd USA/Europe Air Traffic Management R&D Seminar, Napoli, Italy.
- Gillen, D., Jacquillat, A., and Odoni, A. (2016). Airport demand management: The operations research and economics perspectives and potential synergies. *Transportation Research Part A*, 94:495–513.

- Glover, C. and Ball, M. (2013). Stochastic optimization models for ground delay program planning with equity-efficiency tradeoffs. *Transportation Research Part C*, 33:196–202.
- Gopalakrishnan, K. and Balakrishnan, H. (2017). *A Comparative Analysis of Models for Predicting Delays in Air Traffic Networks*. Proc. 12th USA/Europe Air Traffic Management R&D Seminar, Seattle, WA.
- Green, L., Kolesar, P., and Svoronos, A. (1991). Some Effects of Nonstationarity on Multiserver Markovian Queuing Systems. *Operations Research*, 39(3):502–511.
- Green, L., Kolesar, P., and Whitt, W. (2007). Coping with time-varying demand when setting staffing requirements for a service system. *Production and Operations Management*, 16(1):13–29.
- Gross, D. and Harris, C. (1998). *Fundamentals of Queueing Theory*. John Wiley & Sons, New York.
- Grunewald, E. (2016). Incentive-based Slot Allocation for Airports. *Transportation Research Procedia*, 14:3761–3770.
- Guadagni, G., Ndreca, S., and Scoppola, B. (2011). Queuing systems with pre-scheduled random arrivals. *Mathematical Methods of Operations Research*, 73(1):1–18.
- Gupta, G., Malik, W., and Jung, Y. (2011). *Effect of uncertainty on deterministic runway scheduling*. Proc. 11th Aviation Technology, Integration and Operations Conference, Reston, VA.
- Gupta, S. (2010). *Transient Analysis of $D(t)/M(t)/1$ Queuing System with Applications to Computing Airport Delays*. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Gupta, S. and Bertsimas, D. (2011). *Multistage Air Traffic Flow Management under Capacity Uncertainty: A Robust and Adaptive Optimization Approach*. Proc. 51st AGIFORS Annual Symposium and Study Group Meeting, Antalya, Turkey.
- Gwiggner, C. and Nagaoka, S. (2014). Data and queueing analysis of a Japanese air-traffic flow. *European Journal of Operational Research*, 235:265–275.
- Haddad, R., Carlier, J., and Moukrim, A. (2008). A new combinatorial approach for coordinating aerial conflicts given uncertainties regarding aircraft speeds. *International Journal of Production Economics*, 112.
- Hall, W. (1999). *Efficient capacity allocation in a collaborative air transportation system*. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Hanley, Z. (2015). *Delay Characterization and Prediction in Major U.S. Airline Networks*. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Hansen, M., Nikoleris, T., Lovell, D., Vlachou, K., and Odoni, A. (2009). *Use of Queueing Models to Estimate Delay Savings from 4D Trajectory Precision*. Proc. 8th USA/Europe Air Traffic Management R&D Seminar, Napa, CA.
- Hassounah, M. and Steuart, G. (1993). Demand for aircraft gates. *Transportation Research Record*, 1423:26–33.
- Hebert, J. and Dietz, D. (1997). Modeling and Analysis of an Airport Departure Process. *Journal of Aircraft*, 34(1):43–47.
- Heidt, A., Helmke, H., Kapolke, M., Liers, F., and Martin, A. (2016). Robust runway scheduling under uncertain conditions. *Journal of Air Transport Management*, 56:28–37.
- Hengsbach, G. and Odoni, A. (1975). *Time dependent estimates of delays and delay costs at major airports*. Technical report, Dept. of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA.
- Hentzen, D., Kamgarpour, M., Soler, M., and Gonzalez-Arribas, D. (2018). On maximizing safety in stochastic aircraft trajectory planning with uncertain thunderstorm development. *Aerospace Science and Technology*, 79:543–553.

- Heyman, D. (1982). On Ross's Conjectures about Queues with Non-stationary Poisson Arrivals. *Journal of Applied Probability*, 19:245–249.
- Hockaday, S. and Kanafani, A. (1974). A methodology for airport capacity analysis. *Transportation Research*, 8:171–180.
- Horonjeff, R. and McKelvey, F. (1994). *Planning and Design of Airports, 4th edition*. McGraw-Hill.
- Hu, J., Prandini, M., and Sastry, S. (2005). Aircraft conflict prediction in the presence of a spatially correlated wind field. *IEEE Transactions on Intelligent Transportation Systems*, 6(3):326–340.
- Hu, X. and Chen, W. (2005). Receding horizon control for aircraft arrival sequencing and scheduling. *IEEE Transactions on Intelligent Transportation Systems*, 6:189–197.
- Hubbard, H. (1978). Terminal Airspace/Airport Congestion Delays. *Interfaces*, 8(2):1–14.
- Idris, H., Clarke, J.-P., Bhuvan, R., and Kang, L. (2002). Queuing model for taxi-out time estimation. *Air Traffic Control Quarterly*, 10(1):1–22.
- Inniss, T. and Ball, M. (2004). Estimating One-Parameter Airport Arrival Capacity Distributions for Air Traffic Flow Management. *Air Traffic Control Quarterly*, 12(3):223–251.
- International Air Transport Association (2019). European national airspace strategies: Fact sheet. https://www.iata.org/pressroom/facts_figures/fact_sheets/Documents/fact-sheet-european-airspace-strategies.pdf. Accessed on Nov 18, 2019.
- Irvine, R. (2002). A Geometrical Approach to Conflict Probability Estimation. *Air Traffic Control Quarterly*, 10(2):85–113.
- Jacquillat, A. (2012). *A queuing model of airport congestion and policy implications at JFK and EWR*. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Jacquillat, A. and Odoni, A. (2015a). An Integrated Scheduling and Operations Approach to Airport Congestion Mitigation. *Operations Research*, 63(6):1390–1410.
- Jacquillat, A. and Odoni, A. (2015b). Endogenous control of service rates in stochastic and dynamic queuing models of airport congestion. *Transportation Research Part E*, 73:133–151.
- Jacquillat, A. and Odoni, A. (2018). A roadmap toward airport demand and capacity management. *Transportation Research Part A*, 114:168–185.
- Jacquillat, A., Odoni, A., and Webster, M. (2017). Dynamic Control of Runway Configurations and of Arrival and Service Departure Rates at JFK Airport Under Stochastic Queue Conditions. *Transportation Science*, 51(1):155–176.
- Jacquillat, A. and Vaze, V. (2018). Interairline Equity in Airport Scheduling Interventions. *Transportation Science*, 52(4):941–964.
- Jeddi, B., Shortle, J., and Sherry, L. (2006). *Statistics of the Approach Process at Detroit Metropolitan Wayne County Airport*. Proc. 2nd International Conference on Research in Air Transportation, Belgrade, Serbia.
- Jilkov, V., Ledet, J., and Li, X. (2018). Multiple Model Method for Aircraft Conflict Detection and Resolution in Intent and Weather Uncertainty. *IEEE Transactions on Aerospace and Electronic Systems*, 55(2):1004–1020.
- Jilkov, V., Li, X., and Ledet, J. (2014). *Improved estimation of conflict probability for aircraft collision avoidance*. Proc. 17th International Conference on Information Fusion, Salamanca, Spain.
- Joint Planning and Development Office (2010). Concept of Operations for the Next Generation Air Transportation System.

- Jones, J., Lovell, D., and Ball, M. (2015). *Combining Control by CTA and Dynamic Enroute Speed Adjustment to Improve Ground Delay Program Performance*. Proc. 11th USA/Europe Air Traffic Management R&D Seminar, Lisbon, Portugal.
- Jones, J., Lovell, D., and Ball, M. (2018). Stochastic Optimization Models for Transferring Delay Along Flight Trajectories to Reduce Fuel Usage. *Transportation Science*, 52(1):134–149.
- Ju, F., Cai, K., Yang, Y., and Gao, Y. (2015). *A Scenario-based Optimization Approach to Robust Estimation of Airport Capacity*. Proc. 18th International Conference on Intelligent Transportation Systems, Las Palmas, Spain.
- Jung, M. and Lee, E. (1989). Numerical Optimization of a Queueing System by Dynamic Programming. *Journal of Mathematical Analysis and Applications*, 141(1):84–93.
- Kall, P. and Mayer, J. (2011). *Stochastic Linear Programming: Models, Theory and Computation, 2nd edition*. Springer.
- Kantas, N., Lecchini Visintini, A., and Maciejowski, J. (2010). Simulation-based Bayesian optimal design of aircraft trajectories for air traffic management. *International Journal of Adaptive Control and Signal Processing*, 24:882–899.
- Kendall, D. (1953). Stochastic Processes Occurring in the Theory of Queues and their Analysis by the Method of the Imbedded Markov Chain. *Annals of Mathematical Statistics*, 24(3):338–354.
- Khassiba, A., Bastin, F., Cafieri, S., Gendron, B., and Mongeau, M. (2020). Two-Stage Stochastic Mixed-Integer Programming with Chance Constraints for Extended Aircraft Arrival Management. *Transportation Science*, 54(4):897–919.
- Kivestu, P. (1976). *Alternative methods of investigating the time-dependent M/G/K queue*. Master’s thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Kleywegt, A., Shapiro, A., and Homem-de Mello, T. (2002). The Sample Average Approximation Method for Stochastic Discrete Optimization. *SIAM Journal on Optimization*, 12(2):479–502.
- Klooster, J., Del Amo, A., and Manzi, P. (2009). *Controlled Time-of-Arrival Flight Trials*. Proc. 8th USA/Europe Air Traffic Management R&D Seminar, Napa, CA.
- Kohl, N., Larsen, A., Larsen, J., Ross, A., and Tiourine, S. (2007). Airline disruption management - Perspectives, experiences and outlook. *Journal of Air Transport Management*, 13:149–162.
- Koopman, B. (1972). Air-Terminal Queues under Time-Dependent Conditions. *Operations Research*, 20(6):1089–1114.
- Kotnyek, B. and Richetta, O. (2006). Equitable Models for the Stochastic Ground-Holding Problem under Collaborative Decision Making. *Transportation Science*, 40(2):133–146.
- Krozel, J. and Peters, M. (1997). *Strategic conflict detection and resolution for free flight*. Proc. 36th IEEE Conference on Decision and Control, San Diego, CA.
- Kuchar, J. and Yang, L. (2000). A Review of Conflict Detection and Resolution Modeling Methods. *IEEE Transactions on Intelligent Transportation Systems*, 1(4):179–189.
- Lancia, C. and Lulli, G. (2020). Predictive modeling of inbound demand at major European airports with Poisson and Pre-Scheduled Random Arrivals. *European Journal of Operational Research*, 280:179–190.
- Lecchini Visintini, A., Glover, W., Lygeros, J., and Maciejowski, J. (2006). Monte Carlo Optimization for Conflict Resolution in Air Traffic Control. *IEEE Transactions on Intelligent Transportation Systems*, 7(4):470–482.
- Lee, D., Kostiuk, P., Hemm, R., Wingrove, W., and Shapiro, G. (1997). *Estimating the effects of the terminal area productivity program*. NS301R3, Logistics Management Institute, VA.

- Lehoullier, T., Nasri, M., Soumis, F., Desaulniers, G., and Omer, J. (2017). Solving the Air Conflict Resolution Problem Under Uncertainty Using an Iterative Biobjective Mixed Integer Programming Approach. *Transportation Science*, 51(4):1242–1258.
- Li, M. and Ryerson, M. (2019). Reviewing the DATAS of aviation research data: Diversity, availability, tractability, applicability, and sources. *Journal of Air Transport Management*, 75:111–130.
- Lian, G., Zhang, Y., Xing, Z., Luo, Q., and Cheng, S. (2019). A new dynamic pushback control method for reducing fuel-burn costs: Using predicted taxi-out time. *Chinese Journal of Aeronautics*, 32(3):660–673.
- Lim, A., Rodrigues, B., and Zhu, Y. (2005). Airport gate scheduling with time windows. *Artificial Intelligence Review*, 24:5–31.
- Liu, M., Liang, B., Zheng, F., Chu, C., and Chu, F. (2018). *A Two-stage Stochastic Programming Approach for Aircraft Landing Problem*. Proc. 2018 International Conference on Service Systems and Service Management (ICSSSM), Hangzhou, China.
- Liu, P., Hansen, M., and Mukherjee, A. (2008). Scenario-based air traffic flow management: From theory to practice. *Transportation Research Part B*, 42:685–702.
- Liu, W. and Hwang, I. (2014). Probabilistic Aircraft Midair Conflict Resolution Using Stochastic Optimal Control. *IEEE Transactions on Intelligent Transportation Systems*, 15(1):37–46.
- Long, D. and Hasan, S. (2009). *Improved Prediction of Flight Delays Using the LMINET2 System-Wide Simulation Model*. Proc. 9th Aviation Technology, Integration and Operations Conference, Hilton Head, SC.
- Long, D., Lee, D., Johnson, J., Gaier, E., and Kostinuk, P. (1999). *Modeling Air Traffic Management Technologies with a Queuing Network Model of the National Airspace System*. Technical report NASA/CR-1999-208988, NASA Langley Research Center, Hampton, VA.
- Lovell, D., Churchill, A., Odoni, A., Mukherjee, A., and Ball, M. (2007). *Calibrating Aggregate Models of Flight Delays and Cancellation Probabilities at Individual Airports*. Proc. 7th USA/Europe Air Traffic Management R&D Seminar, Barcelona, Spain.
- Lovell, D., Vlachou, K., Rabbani, T., and Bayen, A. (2013). A diffusion approximation to a single airport queue. *Transportation Research Part C*, 33:227–237.
- Lulli, G. and Odoni, A. (2007). The European Air Traffic Flow Management Problem. *Transportation Science*, 41(4):431–443.
- Lygeros, J. and Prandini, M. (2002). *Aircraft and weather models for probabilistic collision avoidance in air traffic control*. Proc. 41st IEEE Conference on Decision and Control, Las Vegas, NV.
- Malone, K. (1995). *Dynamic queuing systems: Behavior and approximations for individual queues and networks*. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Mangoubi, R. and Mathaisel, D. (1985). Optimizing gate assignments at airport terminals. *Transportation Science*, 19:173–188.
- Marianov, V. and Serra, D. (2003). Location models for airline hubs behaving as $M/D/c$ queues. *Computers & Operations Research*, 30:983–1003.
- Matsuno, Y., Andreeva-Mori, A., and Matayoshi, N. (2017). *Effect of Uncertainty on Dynamic Scheduling of Runway Operations*. Proc. 17th Aviation Technology, Integration and Operations Conference, Denver, CO.
- Matsuno, Y., Tsuchiya, T., Wei, J., Hwang, I., and Matayoshi, N. (2015). Stochastic optimal control for aircraft conflict resolution under wind uncertainty. *Aerospace Science and Technology*, 43:77–88.

- McFarlane, P. and Balakrishnan, H. (2016). *Optimal Control of Airport Pushbacks in the Presence of Uncertainties*. Proc. American Control Conference (ACC), Boston, MA.
- Meijers, N. and Hansman, R. (2019). *Data-driven predictive analytics of runway occupancy time for improved capacity at airports*. Report No. ICAT-2019-14, Dept. of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA.
- Mercer, A. (1960). A Queueing Problem in which the Arrival Times of the Customers are Scheduled. *Journal of the Royal Statistical Society, Series B (Methodological)*, 22(1):108–113.
- Moser, I. and Hendtlass, T. (2007). *Solving dynamic single-runway aircraft landing problems with extremal optimization*. Proc. IEEE symposium on computational intelligence in scheduling, Honolulu, HI.
- Mukherjee, A. and Hansen, M. (2007). A Dynamic Stochastic Model for the Single Airport Ground Holding Problem. *Transportation Science*, 41(4):444–456.
- Mukherjee, A. and Hansen, M. (2009). A dynamic rerouting model for air traffic flow management. *Transportation Research Part B*, 43:159–171.
- Mukherjee, A., Lovell, D., Ball, M., Odoni, A., and Zerbib, G. (2005). *Modeling delays and cancellation probabilities to support strategic simulations*. Proc. 6th USA/Europe Air Traffic Management R&D Seminar, Baltimore, MD.
- Munoz, A., Scarlatti, D., and Costas, P. (2018). *Real-time prediction of flight arrival times using surveillance information*. Proc. 12th European Conference on Software Architecture, Madrid, Spain.
- Murca, M. and Muller, C. (2015). Control-based optimization approach for aircraft scheduling in a terminal area with alternative arrival routes. *Transportation Research Part E*, 73:96–113.
- Narciso, M. and Piera, M. (2015). Robust gate assignment procedures from an airport management perspective. *Omega*, 50:82–95.
- Nelson, B. (2013). *Foundations and methods of stochastic simulation: a first course*. Springer Science & Business Media.
- Netjasov, F. (2012). Framework for airspace planning and design based on conflict risk assessment Part 1: Conflict risk assessment model for airspace strategic planning. *Transportation Research Part C*, 24:190–212.
- Newell, G. (1979). Airport Capacity and Delays. *Transportation Science*, 13(3):201–241.
- Ng, K., Lee, C., and Chan, F. (2017). *A robust optimisation approach to the aircraft sequencing and scheduling problem with runway configuration planning*. Proc. 2017 IEEE International Conference on Industrial Engineering and Engineering Management, Singapore.
- Niendorf, M., Kabamba, P., and Girard, A. (2016). Stability Analysis of Runway Schedules. *IEEE Transactions on Intelligent Transportation Systems*, 17(12):3380–3390.
- Nikoleris, T. and Hansen, M. (2012). Queueing models for trajectory-based aircraft operations. *Transportation Science*, 46(4):501–511.
- Nikoleris, T. and Hansen, M. (2015). Effect of Trajectory Prediction and Stochastic Runway Occupancy Times on Aircraft Delays. *Transportation Science*, 50(1):110–119.
- Nilim, A. and El Ghaoui, L. (2004). *Algorithms for Air Traffic Flow Management Under Stochastic Environments*. Proc. American Control Conference (ACC), Boston, MA.

- 1
2
3
4
5
6 Nilim, A., El Ghaoui, L., Duong, V., and Hansen, M. (2001). *Trajectory-Based Air Traffic Management (TB-ATM) Under*
7 *Weather Uncertainty*. Proc. 4th USA/Europe Air Traffic Management R&D Seminar, Santa Fe, NM.
- 8
9 Odoni, A. (1987). *The Flow Management Problem in Air Traffic Control*. In: Odoni, A.R. and Szego, G. (eds). *Flow Control*
10 *of Congested Networks*. Springer-Verlag, Berlin, Germany.
- 11
12 Odoni, A., Morisset, T., Drotleff, W., and Zock, A. (2011). *Benchmarking Airport Airside Performance: FRA vs. EWR*. Proc.
13 9th USA/Europe Air Traffic Management R&D Seminar, Berlin, Germany.
- 14
15 Odoni, A. and Roth, E. (1983). An Empirical Investigation of the Transient Behavior of Stationary Queuing Systems.
16 *Operations Research*, 31(3):432–455.
- 17
18 Omer, J. (2015). A space-discretized mixed-integer linear model for air-conflict resolution with speed and heading maneuvers.
19 *Computers & Operations Research*, 58:75–86.
- 20
21 Paielli, R. (1998). *Empirical test of conflict probability estimation*. Proc. 2nd USA/Europe Air Traffic Management R&D
22 Seminar, Orlando, FL.
- 23
24 Paielli, R. and Erzberger, H. (1997). Conflict probability estimation for free flight. *Journal of Guidance, Control and Dynamics*,
25 20(3):588–596.
- 26
27 Pallottino, L., Feron, E., and Bicchi, A. (2002). Conflict Resolution Problems for Air Traffic Management Systems Solved
28 With Mixed Integer Programming. *IEEE Transactions on Intelligent Transportation Systems*, 3(1):3–11.
- 29
30 Pellegrini, P., Bolic, T., Castelli, L., and Pesenti, R. (2017). SOSTA: An effective model for the Simultaneous Optimisation
31 of airport Slot Allocation. *Transportation Research Part E*, 99:34–53.
- 32
33 Pellegrini, P., Castelli, L., and Pesenti, R. (2012). Secondary trading of airport slots as a combinatorial exchange. *Transporta-*
34 *tion Research Part E*, 48(5):1009–1022.
- 35
36 Pels, E. and Verhoef, E. (2004). The economics of airport congestion pricing. *Journal of Urban Economics*, 55(2):257–277.
- 37
38 Peterson, M., Bertsimas, D., and Odoni, A. (1995). Decomposition algorithms for analyzing transient phenomena in multiclass
39 queueing networks in air transportation. *Operations Research*, 43:995–1011.
- 40
41 Powell, W. (2007). *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley & Sons, New Jersey.
- 42
43 Prandini, M., Hu, J., Lygeros, J., and Sastry, S. (2000). A Probabilistic Approach to Aircraft Conflict Detection. *IEEE*
44 *Transactions on Intelligent Transportation Systems*, 1(4):199–220.
- 45
46 Psaraftis, H. (1978). *A Dynamic Programming Approach to the Aircraft Sequencing Problem*. Technical report R78-4, MIT
47 Flight Transportation Laboratory.
- 48
49 Pujet, N., Delcaire, B., and Feron, E. (1999). *Input-output modeling and control of the departure process of congested airports*.
50 Proc. AIAA Guidance, Navigation and Control Conference, Portland, OR.
- 51
52 Puterman, M. (2005). *Markov Decision Processes - Discrete Stochastic Dynamic Programming, 2nd edition*. Wiley & Sons,
53 New York.
- 54
55 Pyrgiotis, N. (2011). *A Stochastic and Dynamic Model of Delay Propagation Within an Airport Network for Policy Analysis*.
56 Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- 57
58 Pyrgiotis, N., Malone, K., and Odoni, A. (2013). Modelling delay propagation within an airport network. *Transportation*
59 *Research Part C*, 27:60–75.
- 60
61
62
63
64
65

- Pyrgiotis, N. and Odoni, A. (2016). On the Impact of Scheduling Limits: A Case Study at Newark Liberty International Airport. *Transportation Science*, 50(1):150–165.
- Ramanujam, V. and Balakrishnan, H. (2009). *Estimation of Arrival-Departure Capacity Tradeoffs in Multi-Airport Systems*. Proc. 48th IEEE Conference on Decision and Control, Shanghai, China.
- Ratcliffe, S. (1989). Automatic Conflict Detection Logic for Future Air Traffic Control. *The Journal of Navigation*, 42(1):92–106.
- Ravizza, S., Atkin, J., Maathuis, M., and Burke, E. (2013). A combined statistical approach and ground movement model for improving taxi time estimations at airports. *Journal of the Operational Research Society*, 64:1347–1360.
- Rebollo, J. and Balakrishnan, H. (2014). Characterization and prediction of air traffic delays. *Transportation Research Part C*, 44:231–241.
- Rey, D., Rapine, C., Fondacci, R., and El Faouzi, N.-E. (2016). Subliminal Speed Control in Air Traffic Management: Optimization and Simulation. *Transportation Science*, 50(1):240–262.
- Ribeiro, N., Jacquillat, A., and Antunes, A. (2019). A Large-Scale Neighborhood Search Approach to Airport Slot Allocation. *Transportation Science*, 53(6):1772–1797.
- Ribeiro, N., Jacquillat, A., Antunes, A., Odoni, A., and Pita, J. (2018). An optimization approach for airport slot allocation under IATA guidelines. *Transportation Research Part E*, 112:132–156.
- Richetta, O. (1995). Optimal Algorithms and a Remarkably Efficient Heuristic for the Ground-Holding Problem in Air Traffic Control. *Operations Research*, 43(5):758–770.
- Richetta, O. and Odoni, A. (1993). Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control. *Transportation Science*, 27(3):228–238.
- Richetta, O. and Odoni, A. (1994). Dynamic Solution to the Ground-Holding Problem in Air Traffic Control. *Transportation Research Part A*, 28(3):167–185.
- Rolski, T. (1981). Queues with Non-stationary Input Stream: Ross’s Conjecture. *Advances in Applied Probability*, (13):603–618.
- Ross, S. (1978). Average Delay in Queues with Non-stationary Poisson Arrivals. *Journal of Applied Probability*, 15:602–609.
- Ross, S. (1983). *Introduction to Stochastic Dynamic Programming*. Academic Press, New York.
- Rue, R. and Rosenshine, M. (1985). The Application of Semi-Markov Decision Processes to Queueing of Aircraft for Landing at an Airport. *Transportation Science*, 19(2):154–172.
- Saraf, A., Hunter, G., Ramamoorthy, K., Cheng, K., Griffin, K., and Nagle, G. (2014). *Robust And Practical Traffic Flow Management Optimization Algorithm For Near-Term Implementation*. Proc. 14th Aviation Technology, Integration and Operations Conference, Atlanta, GA.
- Saraf, A., Ramamoorthy, K., Hunter, G., Nagle, G., and Yu, P. (2012). *Robust Air Traffic Flow Management: An Optimization-Based Approach*. Proc. 12th Aviation Technology, Integration and Operations Conference, Indianapolis, IN.
- Schwarz, J., Selinka, G., and Stolletz, R. (2016). Performance analysis of time-dependent queueing systems: Survey and classification. *Omega*, 63:170–189.
- Seker, M. and Noyan, N. (2012). Stochastic optimization models for the airport gate assignment problem. *Transportation Research Part E*, 48:438–459.

- Selim Akturk, M., Atamturk, A., and Gurel, S. (2014). Aircraft Rescheduling with Cruise Speed Control. *Operations Research*, 62(4):829–845.
- Sherali, H., Staats, R., and Trani, A. (2003). An Airspace Planning and Collaborative Decision-Making Model: Part I—Probabilistic Conflicts, Workload, and Equity Considerations. *Transportation Science*, 37(4):434–456.
- Sherali, H., Staats, R., and Trani, A. (2006). An Airspace Planning and Collaborative Decision-Making Model: Part II—Cost Model, Data Considerations, and Computations. *Transportation Science*, 40(2):147–164.
- Shone, R., Glazebrook, K., and Zografos, K. (2018). *Stochastic Modelling of Aircraft Queues: A Review*. Proceedings of the OR60 conference, Lancaster, UK.
- Shone, R., Glazebrook, K., and Zografos, K. (2019). Resource allocation in congested queueing systems with time-varying demand: An application to airport operations. *European Journal of Operational Research*, 276(2):566–581.
- Shortle, J., Gross, D., and Mark, B. (2003). *Efficient Simulation of the National Airspace System*. Proc. 35th Winter Simulation Conference, New Orleans, LA.
- Shumsky, R. (1995). *Dynamic statistical models for the prediction of aircraft take-off times*. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Simaiakis, I. (2013). *Analysis, modeling and control of the airport departure process*. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Simaiakis, I. and Balakrishnan, H. (2009). *Queueing models of airport departure processes for emissions reduction*. Proc. AIAA Guidance, Navigation and Control Conference, Chicago, IL.
- Simaiakis, I. and Balakrishnan, H. (2013). *On the Probabilistic Modeling of Runway Inter-departure Times*. Proc. Interdisciplinary Science for Innovative Air Traffic Management (ISIATM) International Conference, Toulouse, France.
- Simaiakis, I. and Balakrishnan, H. (2016). A Queueing Model of the Airport Departure Process. *Transportation Science*, 50(1):94–109.
- Simaiakis, I., Sandberg, M., and Balakrishnan, H. (2014). Dynamic Control of Airport Departures: Algorithm Development and Field Evaluation. *IEEE Transactions on Intelligent Transportation Systems*, 15(1):285–295.
- Skaltsas, G. (2011). *Analysis of Airline Schedule Padding on U.S. Domestic Routes*. Master’s thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Solak, S., Solveling, G., Clarke, J.-P., and Johnson, E. (2018). Stochastic Runway Scheduling. *Transportation Science*, 52(4):917–940.
- Solvering, G., Solak, S., Clarke, J.-P., and Johnson, E. (2011). Runway operations optimization in the presence of uncertainties. *Journal of Guidance, Control and Dynamics*, 34(5):1373–1382.
- Sridhar, B., Grabbe, S., and Mukherjee, A. (2008). Modeling and Optimization in Traffic Flow Management. *Proceedings of the IEEE*, 96(12):2060–2080.
- Stamatopoulos, M., Zografos, K., and Odoni, A. (2004). A decision support system for airport strategic planning. *Transportation Research Part C*, 12:91–117.
- Stolletz, R. (2008). Non-stationary delay analysis of runway systems. *OR Spectrum*, 30:191–213.
- Summers, S. and Lygeros, J. (2010). Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem. *Automatica*, 46(12):1951–1961.

- Sutton, R. and Barto, A. (1998). *Reinforcement Learning: An Introduction*. The MIT Press.
- Swaroop, P., Zou, B., Ball, M., and Hansen, M. (2012). Do more US airports need slot controls? A welfare based approach to determine slot levels. *Transportation Research Part B*, 46(9):1239–1259.
- Tandale, M., Sengupta, P., Menon, P., Cheng, V., Rosenberger, J., and Subbarao, K. (2008). *Queuing Network Models of the National Air Space System*. Proc. 26th Congress of International Council of the Aeronautical Sciences, Anchorage, AK.
- Tang, J. (2019). Conflict Detection and Resolution for Civil Aviation: A Literature Survey. *IEEE Aerospace and Electronic Systems Magazine*, 34(10):20–24.
- Taylor, C. and Wanke, C. (2013). *Designing Coordinated Initiatives for Strategic Traffic Flow Management*. Proc. 10th USA/Europe Air Traffic Management R&D Seminar, Chicago, IL.
- Taylor, C., Wanke, C., Wan, Y., and Roy, S. (2012). *A Decision Support Tool for Flow Contingency Management*. Proc. AIAA Guidance, Navigation and Control Conference, Minneapolis, MN.
- Terrab, M. and Odoni, A. (1993). Strategic Flow Management for Air Traffic Control. *Operations Research*, 41(1):138–152.
- Tielrooij, M., Borst, C., van Paassen, M., and Mulder, M. (2015). *Predicting Arrival Time Uncertainty from Actual Flight Information*. Proc. 11th USA/Europe Air Traffic Management R&D Seminar, Lisbon, Portugal.
- Tien, S.-L., Taylor, C., Zhou, Y., Wan, Y., and Wanke, C. (2011). *A Route-Based Queuing Network Model for Air Traffic Flow Contingency Management*. Proc. 11th Aviation Technology, Integration and Operations Conference, Virginia Beach, VA.
- Tobaruela, G., Fransen, P., Schuster, W., Ochieng, W., and Majumdar, A. (2014). Air traffic predictability framework - Development, performance evaluation and application. *Journal of Air Transport Management*, 39:48–58.
- Tomlin, C., Pappas, G., and Sastry, S. (1998). Conflict Resolution for Air Traffic Management: A Study in Multiagent Hybrid Systems. *IEEE Transactions on Automatic Control*, 43(4):509–521.
- Tu, Y., Ball, M., and Jank, W. (2008). Estimating Flight Departure Delay Distributions - A Statistical Approach With Long-Term Trend and Short-Term Pattern. *Journal of the American Statistical Association*, 103(481):112–125.
- Vaze, V. and Barnhart, C. (2012). An assessment of the impact of demand management strategies for efficient allocation of airport capacity. *International Journal of Revenue Management*, 6:5–27.
- Vela, A., Salaun, E., Solak, S., Feron, E., Singhose, W., and Clarke, J.-P. (2009). *A two-stage stochastic optimization model for air traffic conflict resolution under wind uncertainty*. Proc. IEEE/AIAA 28th Digital Avionics Systems Conference, Orlando, FL.
- Vossen, T., Hoffman, R., and Mukherjee, A. (2011). *Air traffic flow management*. In: Barnhart, C. and Smith, B. (eds). Quantitative Problem Solving Methods in the Airline Industry: A Modelling Methodology Handbook. Springer, Boston, MA.
- Vranas, P., Bertsimas, D., and Odoni, A. (1994a). Dynamic Ground-Holding Policies for a Network of Airports. *Transportation Science*, 28(4):275–291.
- Vranas, P., Bertsimas, D., and Odoni, A. (1994b). The Multi-Airport Ground-Holding Problem in Air Traffic Control. *Operations Research*, 42(2):249–261.
- Wan, Y., Taylor, C., Wanke, C., Roy, S., and Zhou, Y. (2011). *Dynamic Queuing Network Model for Flow Contingency Management*. Proc. AIAA Guidance, Navigation and Control Conference, Portland, OR.

- Wang, K. and Jacquillat, A. (2020). A Stochastic Integer Programming Approach to Air Traffic Scheduling and Operations. *Operations Research*, 68(5):1375–1402.
- Wang, N., Roongnat, C., Rosenberger, J., Menon, P., Subbarao, K., Sengupta, P., and Tandale, M. (2018). Study of time-dependent queuing models of the national airspace system. *Computers & Industrial Engineering*, 117:108–120.
- Wang, R., Alligier, R., Allignol, C., Barnier, N., Durand, N., and Gondran, A. (2020). Cooperation of combinatorial solvers for en-route conflict resolution. *Transportation Research Part C*, 114:36–58.
- Wanke, C. (1997). *Using Air-Ground Data Link to Improve Air Traffic Management Decision Support System Performance*. Proc. 1st USA/Europe Air Traffic Management R&D Seminar, Saclay, France.
- Wanke, C., Taylor, C., Masek, T., Roy, S., and Wan, Y. (2012). *Modeling Air Traffic Demand for a Real-Time Queuing Network Model of the National Airspace System*. Proc. AIAA Modeling and Simulation Technologies Conference, Minneapolis, MN.
- Watkins, O. and Lygeros, J. (2003). *Stochastic Reachability for Discrete Time Systems: An Application to Aircraft Collision Avoidance*. Proc. 42nd IEEE Conference on Decision and Control, Maui, HI.
- Welch, J. and Lloyd, R. (2001). *Estimating Airport System Delay Performance*. Proc. 4th USA/Europe Air Traffic Management R&D Seminar, Santa Fe, NM.
- Weld, C., Duarte, M., and Kincaid, R. (2010). A Runway Configuration Management Model with Marginally Decreasing Transition Capacities. *Advances in Operations Research*, 5(5):490–503.
- Willemain, T., Fan, H., and Ma, H. (2004). *Statistical analysis of intervals between projected airport arrivals*. DSES Technical Report 38-04-510, Rensselaer Polytechnic Institute, Troy, NY.
- Winsten, C. (1959). Journal of the Royal Statistical Society, Series B (Methodological). *Geometric Distributions in the Theory of Queues*, 21(1):1–35.
- Xue, M. and Zelinski, S. (2015). *A Stochastic Scheduler for Integrated Arrival, Departure and Surface Operations in Los Angeles*. Proc. 15th Aviation Technology, Integration and Operations Conference, Dallas, TX.
- Yan, S. and Chang, C. (1998). A network model for gate assignment. *Journal of Advanced Transportation*, 32:176–189.
- Yan, S., Shieh, C.-Y., and Chen, M. (2002). A simulation framework for evaluating airport gate assignments. *Transportation Research Part A*, 36:885–898.
- Yan, S. and Tang, C. (2007). A heuristic approach for gate assignments for stochastic flight delays. *European Journal of Operational Research*, 180:547–567.
- Yang, Y. (2018). Practical Method for 4-Dimensional Strategic Air Traffic Management Problem With Convective Weather Uncertainty. *IEEE Transactions on Intelligent Transportation Systems*, 19(6):1697–1708.
- Yang, Y., Zhang, J., Cai, K.-Q., and Prandini, M. (2017). Multi-aircraft Conflict Detection and Resolution Based on Probabilistic Reach Sets. *IEEE Transactions on Control Systems Technology*, 25(1):309–316.
- Yimga, J. and Gorjidoz, J. (2019). Airline schedule padding and consumer choice behavior. *Journal of Air Transport Management*, 78:71–79.
- Zambon, V. (2018). *Stochastic optimization of airport delays controlled by a dynamic programming algorithm: an application to Marco Polo airport in Venice*. Master’s thesis, University of Padova, Italy.
- Zhou, Y., Wan, Y., Roy, S., Taylor, C., and Wanke, C. (2011). *A Stochastic Modeling and Analysis Approach to Strategic Traffic Flow Management under Weather Uncertainty*. Proc. AIAA Guidance, Navigation and Control Conference, Portland, OR.

Zografos, K., Andreatta, G., and Odoni, A. (2013). *Modelling and Managing Airport Performance*. Wiley-Blackwell.

Zografos, K. and Jiang, Y. (2019). A Bi-objective efficiency-fairness model for scheduling slots at congested airports. *Transportation Research Part C*, 102:336–350.

Zografos, K., Madas, M., and Androutsopoulos, K. (2017). Increasing airport capacity utilisation through optimum slot scheduling: review of current developments and identification of future needs. *Journal of Scheduling*, 20(1):3–24.

Zografos, K., Salouras, Y., and Madas, M. (2012). Dealing with the efficient allocation of scarce resources at congested airports. *Transportation Research Part C*, 21:244–256.

Zou, J., Ahmed, S., and Sun, X. (2019). Stochastic dual dynamic integer programming. *Mathematical Programming*, 175:347–374.

Journal Pre-proof