ARTICLE IN PRESS
JID: JARAP [m5G;April 27, 2016;16:47]
Annual Reviews in Control 000 (2016) 1–8
Contents lists available at ScienceDirect
Annual Reviews in Control
journal homepage: www.elsevier.com/locate/arcontrol

Full Length Article

Control and optimization algorithms for air transportation systems

Hamsa Balakrishnan

Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

ARTICLE INFO

Article history:
Received 27 January 2016
Revised 16 February 2016
Accepted 8 March 2016
Available online xxx

Keywords:
Air transportation
Congestion control
Large-scale optimization
Data-driven modeling
Human decision processes

ABSTRACT

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

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

1. Introduction

The air transportation system operated nearly 85 million flights worldwide in 2014, serving 6.7 billion passengers and 102 million metric tons of cargo. The Asia-Pacific region served more than a third of these passengers, while Europe and North America served about a quarter each. Emerging markets in the Middle East are experiencing an annual growth in traffic of more than 10% annually (Airports Council International, 2015). Although there are nearly 42,000 airports worldwide (nearly 20,000 airports in the United States), traffic demand tends to be concentrated at a small number of them: The top 30 airports serve more than one-third of all passengers, while the busiest airports (Chicago O'Hare, Atlanta and Los Angeles) each see more than 700,000 aircraft operations annually (Airports Council International, 2015; Central Intelligence Agency, 2015).

The increasing demand for air traffic operations has further strained this already capacity-limited system, leading to significant congestion, flight delays, and pollution. Domestic flight delays in the US have been estimated to cost airlines over $19 billion and the national economy over $41 billion annually, waste 740 million gallons of jet fuel, and release an additional 7.1 billion kilograms of CO_2 into the earth's atmosphere (Joint Economic Committee, US Senate, 2008). The demand for airspace resources is expected to significantly in the upcoming decades, and to also include operations of autonomous aircraft (Joint Planning & Development Office, 2004; United States Government Accountability Office, 2015). The networked nature of the air transportation system also leads to the propagation of delays from one part of the system to another. To prevent cascading delays and even congestive collapse, there is a need for new analysis techniques and operational strategies for air transportation systems.

The design of algorithms for air transportation, as in the case of most real-world infrastructures, yields a range of multi-objective optimization problems: For example, one would like to improve the efficiency (in terms of reducing total flight delays, fuel burn, delays per passenger, etc.), robustness (that is, minimize the propagation of delays through the system), while still maintaining the safety and security of the system. These objectives are difficult to achieve in practice, due to the challenges posed by the presence of uncertainties, human factors, and competing stakeholder interests. However, it is possible to overcome these challenges by leveraging the increasingly available operational data to build simple yet realistic models, and to use these models to develop and implement scalable control and optimization algorithms to improve system performance.

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

1. Airport congestion control.
2. Large-scale optimization algorithms for air traffic flow management.

E-mail address: hamsa@mit.edu

http://dx.doi.org/10.1016/j.arcontrol.2016.04.019
1367-5788/

Please cite this article as: H. Balakrishnan, Control and optimization algorithms for air transportation systems, Annual Reviews in Control (2016), http://dx.doi.org/10.1016/j.arcontrol.2016.04.019

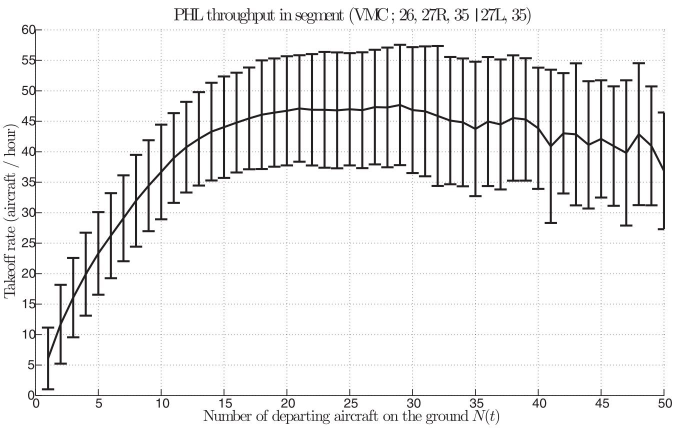

Fig. 1. Average take-off rate as a function of the number of departing aircraft on the ground at PHL. The error bars represent the standard deviation of the take-off rate (Simaiakis & Balakrishnan, 2010).

3. Learning models of air traffic controller decision processes and the associated utility functions.

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

2. Airport congestion control

Taxiing aircraft consume fuel, and emit pollutants such as Carbon Dioxide, Hydrocarbons, Nitrogen Oxides, Sulfur Oxides and Particulate Matter that impact the local air quality at airports (Ashok, Dedoussi, Yim, Balakrishnan, & Barrett, 2014; Carslaw, Beevers, Ropkins, & Bell, 2006; Miracolo et al., 2011; Yu, Cheung, Cheung, & Henry, 2004). Although fuel burn and emissions are approximately proportional to the taxi times of aircraft, other factors such as the throttle settings, number of engines that are powered, and pilot and airline decisions regarding engine shutdowns during delays also influence them (Simaiakis & Balakrishnan, 2010). Domestic flights in the United States emit about 6 million metric tonnes of CO_2, 45,000 tonnes of CO, 8000 tonnes of NOx, and 4000 tonnes of HC taxiing out for takeoff; almost half of these emissions are at the 20 most congested airports in the country (Simaiakis, Khadilkar, Balakrishnan, Reynolds, & Hansman, 2014). Aircraft in Europe have been estimated to spend 10–30% of their flight time taxiing (Cros & Frings, 2008). Data also show that 20% of delays at major US airports occur not due to bad weather, but due to high traffic volume (Federal Aviation Administration (FAA), 2015). Better congestion management at airports has the potential to mitigate these impacts.

2.1. Impacts of airport congestion

Pujet et al. analyzed surface congestion by considering the take-off rate of an airport as a function of the number of aircraft taxiing out (Pujet, Delcaire, & Feron, 2000). Fig. 1 shows a similar analysis for Philadelphia International Airport (PHL) in 2007, for one runway configuration (set of active runways at the time), under visual meteorological conditions (VMC) (Simaiakis & Balakrishnan, 2010).

Fig. 1 illustrates that although the take-off rate increases at first, it saturates once there are approximately 20 departing aircraft on the ground. Any further pushbacks will just lead to congestion, and will not result in an improvement in the takeoff rate. It is also worth noting that for a very high numbers of departures on the ground (more than 30 in Fig. 1), the departure throughput can even *decrease* due to surface gridlock. Similar phenomena have been observed at several major airports in the US, including Boston Logan International Airport (BOS), Newark Liberty International Airport (EWR), New York John F. Kennedy International Airport (JFK), New York La Guardia International Airport (LGA), and Charlotte Douglas International Airport (CLT) (Sandberg, Reynolds, Khadilkar, & Balakrishnan, 2013; Simaiakis, 2012; Simaiakis & Balakrishnan, 2010; 2015). This phenomenon of throughput saturation is also typical of queuing systems, motivating the development of queuing network models of major airports (Jacquillat, 2012; Simaiakis & Balakrishnan, 2015).

2.2. Congestion management strategies

One of the earliest efforts at airport congestion control was the Departure Planner project (Feron et al., 1997). This project proposed the concept of a virtual departure queue, where aircraft would be held (at their gates) until an appropriately determined pushback time. The resultant *N-Control* strategy was a threshold heuristic, where if the total number of departing aircraft on the ground exceeded a certain threshold, N_{ctrl}, any further aircraft requesting pushback were held at their gates until the number of departures on the ground fell below the threshold (Feron et al., 1997; Simaiakis, Khadilkar et al., 2014). Other variants and extensions of this policy have also been studied (Burgain, Feron, Clarke, & Darrasse, 2008; Carr, 2001; Carr, Evans, Feron, & Clarke, 2002; Pujet et al., 2000). Interestingly, a similar heuristic has been known to be deployed by Air Traffic Controllers at BOS during times of extreme congestion (Clewlow & Michalek, 2010). The N-Control policy is similar in spirit to *constant work-in-process* or CONWIP policies that have been proposed for manufacturing systems (Spearman & Zazanis, 1992).

Several other approaches to departure metering have been proposed, including the Ground Metering Program at New York's JFK airport (Nakahara, Reynolds, White, & Dunsky, 2011; Stroiney, Levy, Khadilkar, & Balakrishnan, 2013), the field-tests of the Collaborative Departure Queue Management concept at Memphis (MEM) airport (Brinton, Provan, Lent, Prevost, & Passmore, 2011), the human-in-the-loop simulations of the Spot and Runway Departure Advisor (SARDA) concept at Dallas Fort Worth (DFW) airport (Jung et al., 2011), and the trials of the Departure Manager (DMAN) concept (Böhme, 2005) at Athens International airport (ATH) (Schaper, Tsoukala, Stavrati, & Papadopoulos, 2011). In addition, Mixed Integer Linear Programming (MILP) formulations of surface traffic optimization have been considered, but are generally known to be NP-hard (Balakrishnan & Jung, 2007; Lee, 2014; Lee & Balakrishnan, 2010; Rathinam, Montoya, & Jung, 2008; Smeltink, Soomer, de Waal, & van der Mei, 2005). In practice, these strategies are treated as open-loop policies that are periodically reoptimized. Full-state feedback policies have also been proposed, but have presented practical challenges (Burgain, Pinon, Feron, Clarke, & Mavris, 2009).

2.3. Design and implementation of a congestion control algorithm

While there has been prior research on the optimal control of queuing systems (Crabill, Gross, & Magazine, 1977; Stidham & Weber, 1993), the application of these techniques to airport operations has remained a challenge. In particular, the need to interface with current air traffic control procedures, and the different sources of uncertainty (the variability in departure throughput and the randomness of taxi-out times) pose practical concerns.

2.3.1. Rate control strategies

On-off or event-driven pushback control policies (such as a threshold heuristic) are not desirable in practice, since both air