Available online at www.sciencedirect.com

ScienceDirect

Transportation Research Part B 42 (2008) 685–702

TRANSPORTATION
RESEARCH
PART B

www.elsevier.com/locate/trb

Scenario-based air traffic flow management: From theory to practice

Pei-chen Barry Liu[a], Mark Hansen[a,*], Avijit Mukherjee[b]

[a] Department of Civil and Environmental Engineering, Institute of Transportation Studies,
National Center of Excellence of Aviation Operations Research, University of California, Berkeley, Berkeley, CA 94720, USA
[b] NASA University Affiliated Research Center, University of California, Santa Cruz, Santa Cruz California, 94035, USA

Received 1 July 2006; received in revised form 4 January 2008; accepted 4 January 2008

Abstract

Recent developments in solving the single airport ground holding problem use static or dynamic optimization to manage uncertainty about how airport capacities will evolve. Both static and dynamic models involve the use of scenarios that depict different possible capacity evolutions. Dynamic models also require scenario trees featuring branch points where previously similar capacity profiles become distinct. In this paper, we present methodologies for generating and using scenario trees from empirical data and examine the performance of scenario-based models in a real-world setting. We find that most US airports have capacity profiles that can be classified into a small number of nominal scenarios, and for a number of airports these scenarios can be naturally combined into scenario trees. The costs incurred from applying scenario-based optimization, either static or dynamic, to these airports is considerably higher than the "theoretical" optimization results suggest because actual capacities vary around the nominal values assumed in the optimization, and because of uncertainty in navigating scenario trees that the idealized models ignore. Methods for tuning capacity scenarios and scenario trees to mitigate these problems are explored.

Keywords: Air traffic flow management; Ground delay program; Ground holding problem; Capacity scenarios; Airport capacity; Stochastic optimization; Dynamic optimization

1. Introduction

Ground Delay Programs (GDPs) are implemented to control air traffic volume inbound to airports where the projected traffic demand is expected to exceed the airport capacity, or acceptance rate, for an extended period of time. Under a GDP, the capacity-demand imbalance is managed through delaying some flights prior to their departure. There has been considerable research on how to allocate ground delays efficiently and

---

[*] Corresponding author. Tel.: +1 510 642 2880; fax: +1 510 642 5687.
E-mail addresses: barryliu@cal.berkeley.edu (P.-c.B. Liu), hansen@ce.berkeley.edu (M. Hansen), avijit@ucsc.edu (A. Mukherjee).

0191-2615/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.trb.2008.01.002

equitably. The Single Airport Ground Holding Problem (SAGHP) has been addressed by deterministic models in the seminal studies (Odoni, 1987) followed by models that consider the uncertainty in airport arrival capacity using stochastic integer programming under static (Ball et al., 2003) and partially dynamic (Richetta and Odoni, 1994) settings. In the Ball et al. (2003) model, ground delays are assigned at the beginning of the planning horizon and not revised later. In the Richetta and Odoni (1994) model ground delays are assigned to groups of flights as their scheduled departure times approach, according to the latest information on capacity conditions at the destination airport. However, once the ground delays are assigned to flights, they are not revised later even for flights that have yet to depart. In a more recent development, the Mukherjee–Hansen model (Mukherjee, 2004; Mukherjee and Hansen, 2007), allows for the revision of ground delays for flights that have not yet departed in response to updated information. This allows for ‘‘wait-and-see” strategies in which flights are held in anticipation of upcoming information.

The Ball et al., Richetta–Odoni, and Mukherjee–Hansen models are all based on the assumption that arrival capacity evolution can be modeled using a limited number of scenarios. The latter two models also assume that scenarios form trees with branch points over the course of the day when previously similar scenarios become distinct. The scenarios reported in the literature were constructed for purposes of illustration. This begs the questions of whether and how scenario-based models can be implemented for real-world airports, and how the benefits of such implementations compare to those suggested by the idealized applications reported to date.

This paper addresses these questions. First, we propose and demonstrate methodologies that can be used to determine airport-specific scenario trees at any airport. Unlike previous studies (Inniss and Ball, 2004) on arrival capacity distributions in the literature, we provide a methodology that does not restrict the patterns by a few pre-postulated forms and organizes the arrival capacity profiles into a probabilistic scenario tree. Additionally, we propose methods for ‘‘navigating” through a scenario tree so that at a given time we know (or can intelligently guess) which branch of the tree we are on.

The ultimate value of these methodologies, and of the models they support, depends upon the benefits they provide in a real-world setting. In this paper, we assess these benefits by simulating what would happen if the air traffic flow management strategies obtained from static and dynamic models are used in the face of actual capacity profiles, which do not match the idealized capacity scenarios perfectly. We compare the results of static (Ball et al.) and dynamic (Mukherjee–Hansen) optimization models in this regard, and also report on experiments using different methods for applying the dynamic model.

The remainder of this paper is organized as follows. In the next section, we introduce terminology and identify the challenges in implementing scenario-based air traffic management in a real-world setting. Next, we present methods for developing scenario trees from data, and illustrate these methods using data from several major US airports. In Section 5, we introduce two scenario-based stochastic optimization models – the static one of Ball et al. and the dynamic Mukherjee–Hansen model – and apply them to the same airports discussed in Section 3. Next we discuss challenges to implementing optimal solutions, methods for overcoming these challenges, and how these challenges and their solutions affect model performance. Section 7 offers conclusions.

2. Using scenario trees in air traffic management

We first define the related concepts of capacity profile, capacity scenario, and scenario tree, that are the focus of the subsequent discussion. We define a *capacity profile* as a time series of capacity values for an airport over a day or some part of a day. Capacity profiles are derived from airport acceptance rates (AARs), which specify the number of arrivals an airport can accommodate over a 15-min period. Thus a capacity profile is a time series of AARs for an airport. We use the term *capacity scenario* to refer to a representative profile that is similar, but not identical, to a set of profiles. If profiles are individuals, then scenarios are species. The process of developing scenarios from profiles will be discussed below. Finally, a *scenario tree* is a set of capacity scenarios with associated probabilities of realization that is organized into a tree where a branching point represents the time when initially similar scenarios evolve into distinct scenarios.

Our aim is to develop practical methods for creating and using scenarios trees for air traffic management. To do this, we must address four specific challenges: