Transportation Research Part B 43 (2009) 466–475

Contents lists available at ScienceDirect

Transportation Research Part B

journal homepage: www.elsevier.com/locate/trb

Improving the objective function of the fleet assignment problem

Jonathan Dumas *, Fati Aithnard, François Soumis

GERAD and Département de Mathématiques et de Génie Industriel, École Polytechnique de Montréal, C.P. 6079, Succ. Centre-ville, Montréal, QC, Canada H3C 3A7

ARTICLE INFO

Article history:
Received 20 December 2007
Received in revised form 18 August 2008
Accepted 25 August 2008

Keywords:
Airline transport optimization
Fleet assignment
Passenger flow
Revenue estimation

ABSTRACT

Most fleet assignment problem (FAP) formulations use a leg-based estimation of revenue loss to derive the passenger revenue component of their objective function. This neglects the leg interdependency of revenues, caused by multileg itineraries. We tackle this problem by modifying the objective function using information provided by a passenger flow model devised by two of the authors. It models spill and recapture between itineraries, accounts for the leg interdependency of revenues and does not control passenger flow to the airline company's advantage. We iteratively improve the FAP's objective function by alternately generating fleet assignments and analyzing them with a modified version of the passenger flow model. We have tested this process on a large-scale network made up of Air Canada data with various demand levels and distributions. Most of the profit improvement occurs in the first few iterations, and the objective function adjustment takes on average less than half the FAP resolution time.

1. Introduction

The fleet assignment problem (FAP) in air transportation consists in assigning fleet types to flights of a given flying schedule while respecting feasibility constraints and maximizing expected profit.

Expected operational costs can be quite accurately expressed as a linear function of the FAP decision variables. The same cannot be said of expected revenues. Future revenues depend mainly on future travel demand. A particular fleet assignment (FA) then affects revenues through capacity limitations. Hence, when one expresses the expected profit as a linear function of the FAP decision variables, it is implicit that it is based on demand forecasts and that this linear function is only an approximation, a tool whose efficiency is to be judged on the profitability of the FAs found by the FAP solver.

Standard FAP formulations use a *leg-based* estimation of revenues (see Section 2.2 that ignores both the dependence between passenger flows on flight legs that share multileg itineraries (*network effects*) and the recapture. Several researchers have proposed strategies to make up for these shortcomings. In Jacobs et al. (1999) (see Smith (2004) for a detailed exposition and an improved technique), Benders decomposition is used to integrate the FAP with an origin-destination revenue management model. The upper bound on revenue in the relaxed master problem, a piecewise linear, concave function, is gradually improved until it is in acceptable agreement with the revenue evaluation of the revenue management model. The FAP then is solved as a mixed integer program. In Kliewer (2000), the author integrates a deterministic passenger flow model to the FAP and uses simulated annealing to solve it. The algorithm is used as part of a much broader integrated strategy for airline operations planning, described in Kliewer et al. (2002) and Weber et al. (2003). In Farkas (1995), network effects are accounted for in a FAP formulation through the use of passenger flow decision variables, assuming airline control

* Corresponding author. Tel.: +1 514 340 6053x6046; fax: +1 514 340 5665.
E-mail addresses: jonathan.dumas@gerad.ca (J. Dumas), Fati.Aithnard@gerad.ca (F. Aithnard), francois.soumis@gerad.ca (F. Soumis).

0191-2615/$ - see front matter © 2008 Elsevier Ltd. All rights reserved.
doi:10.1016/j.trb.2008.08.005

over passenger flows. This approach is pursued in Barnhart et al. (2002), where recapture is modelled as well, and where demand is deterministic.

We present here an iterative improvement scheme for the linear objective function of a standard FAP formulation. It relies on the computation, at each iteration, of a *post-FA revenue loss vector*, defined in Section 2.3. The scheme is conceptually simple but requires tens of thousands of calls to a flow model algorithm to generate the next objective function. In fact, the main work behind this paper has been to test numerous modifications, or localizations, of our passenger flow model to find a right balance between accuracy and computing speed. The solution we retain is described in Section 3.2. We use a standard leg-based estimation of revenue losses to generate the first objective function, and the initial FA it yields is our basis of comparison for the subsequent, improved, FAs.

Our approach is novel in that it makes the FAP objective function carry the desirable characteristics of the underlying passenger flow model, described by two of the authors in Dumas and Soumis (2008). This passenger flow model respects the stochastic nature of the demand and the temporality of the booking process. It models network effects, spill and recapture, and, importantly, it does not treat passenger flows as decision variables whose values are set to maximize revenues according to a mathematical program.

2. General scheme

2.1. Notation

Let us first list and explain the notation we use.

- $\mathscr{N}$ is an airline network, typically covering a week of service;
- $\mathscr{L}$ is the set of its flight legs;
- T_l is the set of admissible fleet types for leg l;
- t^∞ is an artificial fleet type of infinite capacity;
- $TFA \subseteq \{0,1\}^{\sum_{l\in\mathscr{L}}|T_l|}$, the set of *tentative fleet assignments* for $\mathscr{N}$, is the set of 0–1 vectors $\mathbf{X}$ having exactly one nonzero entry $\mathbf{X}_{l,t}$ for each leg $l \in \mathscr{L}$;
- $FA \subseteq TFA$ is the set of *legal fleet assignments* for $\mathscr{N}$.

We work with the basic FAP formulation as a multicommodity network flow problem described in Hane et al. (1995). By *legal* FA, we mean one respecting the feasibility constraints: cover (exactly one fleet type must be assigned to each flight leg); plane count (for each fleet type, the number of airplanes to be used at any moment must not exceed the airline's fleet size); and balance (flow conservation in the subnetworks induced by each fleet type).

If $\mathbf{X} \in TFA$, then $\mathbf{X}_{l,t} = 1$ means that $\mathbf{X}$ assigns fleet type t to leg l. We write $\mathbf{X}(l) = t$ to mean that $\mathbf{X}_{l,t} = 1$.

Thus, we can write the FAP quite compactly as

$$\begin{aligned} \min \quad & \sum_{l\in\mathscr{L}, t\in T_l} \mathbf{X}_{l,t}(C_{l,t} + L_{l,t}) \\ \text{s.t.} \quad & \mathbf{X} \in FA, \end{aligned}$$

where $C_{l,t}$ is the cost of assigning type t to leg l, and $L_{l,t}$ is the estimated loss of revenue caused by the assignment of type t to leg l.

For $\mathbf{X} \in TFA$, we write

$$\mathbf{X} \oplus [l, t]$$

to denote the tentative fleet assignment that assigns type t to leg l and type $\mathbf{X}(l')$ to any other leg l'.

2.2. Leg-based estimation of revenue losses

Let us subdivide each flight leg of $\mathscr{N}$ into *arcs*, each arc corresponding to one of several aggregate fare classes. For any fleet type t, each arc a has a seating capacity $cap_{a,t}$, and these add up to the seating capacity of type t. We let $\mathscr{A}$ be the set of arcs of $\mathscr{N}$ and we write $a \lhd l$ to mean that a is part of leg l.

We call *itinerary* a sequence of incident arcs of $\mathscr{N}$ of the same fare class. Let $\mathscr{I}$ be the set of itineraries of $\mathscr{N}$ made available to customers. If $i \in \mathscr{I}$, we write $a \in i$ to mean that a is an arc of the itinerary i.

Demands for itineraries vary from week to week and we accordingly consider them as random variables. Their distributions are generally modelled as normal truncated at zero, or gamma, for small demands (Swan, 2002). As in Dumas and Soumis (2008), we assume that, for all $i \in \mathscr{I}$, the forecast demand is provided as a truncated normal random variable D_i of expectation d_i, with coefficient of variation 0.3 when $d_i \geqslant 5$ and 0.5 otherwise. The forecasted demand for arc a is $D_a = \sum_{i:a\in i} D_i$, with expectation d_a.

The revenue loss vector L^0 that we want to improve upon is computed with the following standard leg-based estimation of spilled revenue: