Computers & Operations Research 45 (2014) 68–86

Contents lists available at ScienceDirect

Computers & Operations Research

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

An integrated scenario-based approach for robust aircraft routing, crew pairing and re-timing

Michelle Dunbar [a,*,1], Gary Froyland [b], Cheng-Lung Wu [c]

[a] SMART Infrastructure Facility, University of Wollongong, Sydney NSW 2522, Australia
[b] School of Mathematics and Statistics, University of New South Wales, Sydney NSW 2052, Australia
[c] School of Aviation, University of New South Wales, Sydney NSW 2052, Australia

ARTICLE INFO

Available online 17 December 2013

Keywords:
Robust airline scheduling
Delay propagation
Airline schedule optimization
Schedule re-timing

ABSTRACT

For reasons of tractability, the airline scheduling problem has traditionally been sequentially decomposed into various stages (e.g. schedule generation, fleet assignment, aircraft routing, and crew pairing), with the decisions from one stage imposed upon the decision-making process in subsequent stages. Whilst this approach greatly simplifies the solution process, it unfortunately fails to capture many dependencies between the various stages, most notably between those of aircraft routing and crew pairing, and how these dependencies affect the propagation of delays through the flight network. In Dunbar et al. (2012) [9] we introduced a new algorithm to accurately calculate and minimize the cost of propagated delay, in a framework that integrates aircraft routing and crew pairing. In this paper we extend the approach of Dunbar et al. (2012) [9] by proposing two new algorithms that achieve further improvements in delay propagation reduction via the incorporation of stochastic delay information. We additionally propose a heuristic, used in conjunction with these two approaches, capable of re-timing an incumbent aircraft and crew schedule to further minimize the cost of delay propagation. These algorithms provide promising results when applied to a real-world airline network and motivate our final integrated aircraft routing, crew pairing and re-timing approach which provides a substantially significant reduction in delay propagation.

1. Introduction

1.1. The airline scheduling problem

The airline scheduling problem involves the construction of timetables for an airline's major resources, namely aircraft and crew. Traditionally, this has been undertaken with a view towards maximizing an airline's overall profit, often with limited consideration given to the stability of such a schedule, or indeed its operational robustness. Such an approach has a tendency to generate schedules that are highly brittle, performing poorly in practice as delays propagate rapidly throughout the network. The Bureau of Transportation Statistics [17] states that approximately 21.5% of all flight legs in the U.S. between the months of July 2010 and July 2011 were delayed more than 15 min – with late arrivals and cancellations together accounting for approximately half of this delay.

In recent years, this has resulted in an ever-increasing discrepancy between planned costs and realized operational costs. As aircraft networks continue to grow, this trend is set to continue, with AhmadBeygi et al. [1] reporting that in 2006, it was estimated that the US airline industry experienced a total of 116.5 million minutes of delay; translating into a $7.7 billion increase in operating costs. Such large discrepancies have prompted airline schedule planners to shift their focus from maximizing profit to maximizing expected profits under uncertainty, by including various types of costs arising from unplanned events.

The airline scheduling problem in its entirety is very complex. The vast number of rules and regulations associated with airports, aircraft, and crew, combined with the global reach of air traffic networks, requires the problem to be broken into manageable pieces to maintain some degree of tractability. Consequently, the traditional airline scheduling problem is typically decomposed into four stages, with the output of one stage used as the input for the subsequent stage(s). The very first stage is known as the *schedule generation* problem. In this step, an airline seeks to construct a schedule of flights where each flight is specified by an “origin, destination, departure date, time and duration” Weide et al. (2010) [21]. The origin

---

[*] Corresponding author.
E-mail address: mdunbar@uow.edu.au (M. Dunbar).
[1] A significant proportion of this research was completed while the author was with the University of New South Wales.

0305-0548/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.cor.2013.12.003

and destination of each flight leg (known as an O–D pair), and additionally the frequency with which they are flown, are determined by the market demand for such pairs and availability of aircraft resources. The second stage, known as *fleet assignment* assigns a particular aircraft type (or fleet) to each flight leg, to appropriately match the size of the aircraft to the intended range (e.g. long-haul vs. domestic) and the expected number of passengers. Typically, the objective is to maximize profit via the minimization of operating expenses and number of spilled passengers.

The third stage, known as *aircraft routing*, is performed separately for each specific fleet type to obtain a minimal cost assignment of aircraft to flights that ensures each flight is covered exactly once by exactly one aircraft. The aircraft routing problem is often further decomposed into two sub-problems: aircraft routing generation and tail assignment. If performed sequentially by these two sub-problems, aircraft routing generation is usually modelled as a feasibility problem that produces generic maintenance feasible routings. As the day of operations draws closer, generic aircraft routings are then assigned for specific aircraft, i.e. tails, for operations. Finally the last stage, known as *crew pairing*, is also performed separately for each fleet type, as flight crew pilots are only certified for one fleet type at any given moment. The objective of crew pairing is to find a minimal cost assignment of crew to flights. A set of crew pairings are constructed that satisfy safety regulations (such as the 8-in-24 rule),[2] and ensure each flight is covered exactly once by exactly one crew group. Feasible pairings are then assigned to airline crew members as part of the crew rostering problem.

1.2. *Integrated re-timing approaches*

Recognizing that schedule generation plays a key role in determining the feasibility of subsequent aircraft and crew assignments, a number of authors have attempted to combine (an approximation of) schedule generation with fleet assignment, aircraft routing and crew pairing. Additionally, various authors have attempted to incorporate extra flexibility (and potentially improve operational robustness) via the introduction of *time windows*. Time windows allow the departure time to fall anywhere within a discretized window, usually extending 10–15 min either side of the originally scheduled departure time. This is achieved through the introduction of additional flight copies; each corresponding to a choice of possible departure times within the discretized time window, along with corresponding connection arcs; see Fig. 1 below.

Desaulniers et al. (1997) [6] introduce time windows on flight departures for the fleet assignment problem. The problem is modelled as a multi-commodity flow in which extra time variables are introduced. The authors solve this problem using branch-and-bound and column generation, in which the column generator is a time-constrained shortest path problem. Rexing et al. (2000) [18] incorporate time windows within the fleet assignment problem, and the time windows are discretized into 5 (and 1) min intervals. Klabjan et al. (2002) [11] address the problem of airline crew scheduling, incorporating time windows to allow more flexibility in the crew-pairing solution. Lan et al. (2006) [13] attempt to incorporate robustness into the schedule by estimating delays for each flight leg and minimizing expected delay. Lan et al. also use time windows (referred to as re-timing) to address the issue of reducing missed connections for passengers. As in the above, they introduce flight copies, and estimate the number of disrupted passengers for each possible connection using a connection-based model, leaving the fleeting and routing solutions unchanged. Samardi (2004) [19] extends the passenger connection model of Lan (2003) [12] and presents an integrated flight departure, re-timing and aircraft routing model that aims to minimize the expected number of misconnecting passengers. Their model attempts to provide any potential misconnecting passengers with alternative recovery options. Lohatepanont and Barnhart (2004) [14] extend the itinerary-based fleet assignment model (IFAM) of Barnhart et al. (2002) [4] to determine market service frequency, departure times and fleet assignments simultaneously. The authors make use of a set of flight legs that may be categorized as mandatory or optional, and assess the worth of a particular itinerary and re-adjust flight leg demand if a particular itinerary is removed, or if the schedule is altered. Belanger et al. (2006) [5] present an integrated model for fleet assignment with time windows for which they assume the schedule is periodic. The authors penalize short connections between flights and make use of profit estimations that integrate and capture both departure time and aircraft type; resulting in a potentially more profitable matching of fleet type with expected passenger demand. Klabjan et al. (2002) [11] partially integrate aircraft routing with crew scheduling. Solving the problem sequentially, the authors add plane-count constraints to the crew scheduling model to obtain a feasible aircraft routing problem. The authors also include time windows to allow more flexibility within the crew scheduling problem. More recently, Weide (2009) [20] proposed a model for the robust and integrated aircraft routing and crew pairing problem with time windows, employing two different methods for its solution. The model partially integrates scheduling decisions via the inclusion of departure time windows and seeks to achieve robustness by penalising aircraft changes for which connection time is less than a specified restricted time.

1.3. *Outline of this paper*

This paper is set out as follows. In Section 2 we introduce key notation and equations that are used extensively throughout this paper. In Section 3 we outline a re-timing heuristic that improves upon the shortcomings of the existing re-timing approaches mentioned above via the *simultaneous* re-timing of aircraft and crew and in such a way as to minimize delay propagation between aircraft and crew. In Section 4 we extend the integrated aircraft routing and crew pairing model of Dunbar et al. [9] by proposing two alternative approaches for incorporating delay *scenarios* within the aircraft routing and crew pairing problems and outline an approach for integrating this with the improved heuristic. This inclusion of multiple delay scenarios allows an airline to incorporate historical primary delay information into the model in a more meaningful way, rather than simply making use of expected delays. Finally in Section 5 we improve upon these algorithms further by including scheduling decisions within the aircraft routing subproblem to provide even better quality solutions.

2. Preliminaries

In this paper we make use of a number of equations and algorithms proposed in Dunbar et al. (2012) [9]. For a complete overview of these algorithms and additional accompanying explanations, the reader is referred to [9]. We now introduce some key notation used

---

[2] The 8-in-24 rule is imposed by the FAA, and requires that crew be given additional rest should the total flying time of a pairing exceed 8 h in a 24 h period. See [3] for further details.