Transportation Research Part B 113 (2018) 70–90

Contents lists available at ScienceDirect

Transportation Research Part B

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

A column generation-based heuristic for aircraft recovery problem with airport capacity constraints and maintenance flexibility

Zhe Liang[a], Fan Xiao[a], Xiongwen Qian[b,a,*], Lei Zhou[a], Xianfei Jin[c], Xuehua Lu[c], Sureshan Karichery[c]

[a] School of Economics & Management, Tongji University, Shanghai 200092, China
[b] School of Mechanical & Aerospace Engineering, Nanyang Technological University, 639798, Singapore
[c] Sabre Airline Solutions, 3150 Sabre Drive, Southlake, TX 76092, USA

ARTICLE INFO

Article history:
Received 17 November 2017
Revised 20 March 2018
Accepted 7 May 2018

Keywords:
Aircraft recovery problem
Disruptions management
Column generation

ABSTRACT

We consider the aircraft recovery problem (ARP) with airport capacity constraints and maintenance flexibility. The problem is to re-schedule flights and re-assign aircraft in real time with minimized recovery cost for airlines after disruptions occur. In most published studies, airport capacity and flexible maintenance are not considered simultaneously via an optimization approach. To bridge this gap, we propose a column generation heuristic to solve the problem. The framework consists of a master problem for selecting routes for aircraft and subproblems for generating routes. Airport capacity is explicitly considered in the master problem and swappable planned maintenances can be incorporated in the subproblem. Instead of discrete delay models which are widely adopted in much of the existing literature, in this work flight delays are continuous and optimized accurately in the subproblems. The continuous-delay model can improve the accuracy of the optimized recovery cost by up to 37.74%. The computational study based on real-world problems shows that the master problem gives very tight linear relaxation with small, often zero, optimality gaps. Large-scale problems can be solved within 6 min and the run time can be further shortened by parallelizing subproblems on more powerful hardware. In addition, from a managerial point of view, computational experiments reveal that swapping planned maintenances may bring a considerable reduction in recovery cost by about 20% and 60%, depending on specific problem instances. Furthermore, the decreasing marginal value of airport slot quota is found by computational experiments.

1. Introduction

1.1. Background and literature review

Disruptions have a large financial impact on the airline industry. In 2016, the U.S. Passenger Carrier Delay Costs were estimated to be $62.55 per minute (Airlines for America, 2017) and the total U.S. flight arrival delays amounted to over 59

* Corresponding author at: School of Mechanical & Aerospace Engineering, Nanyang Technological University, 639798, Singapore.
E-mail addresses: liangzhe@tongji.edu.cn (Z. Liang), 1631022@tongji.edu.cn (F. Xiao), xqian003@e.ntu.edu.sg (X. Qian), zjzhoulei@tongji.edu.cn (L. Zhou), Xianfei.Jin@sabre.com (X. Jin), Xuehua.Lu@sabre.com (X. Lu), Sureshan.Karichery@sabre.com (S. Karichery).

https://doi.org/10.1016/j.trb.2018.05.007

million minutes (Department of Transportation, 2016). It is critical to develop computational tools for airlines to deal with disruptions and obtain high-quality recovery solutions in real time. The aircraft recovery problem (ARP), as a fundamental part of Airline Disruptions Management, plays a vital role in every airline's daily operation. Although effective methods (Barnhart et al., 1998; Liang and Chaovalitwongse, 2013; Haouari et al., 2013; Liang et al., 2015; Shao et al., 2017) have been proposed to help produce good pre-operational plans for airlines, each day, plans, when implemented, are inevitably subject to unpredictable disruptions that force airlines to make modifications in a timely manner. Disruptions may be caused by airline resource shortages such as aircraft mechanical failures or absence of crew members, or they could be due to airport capacity and air traffic control restrictions such as the quota of available airport departure and arrival slots in adverse weather. When disruptions happen, the Airline Operations Center is responsible for making decisions, re-scheduling airline resources including aircraft, flights and crews, and re-accommodating passengers with the objective of restoring the airline's operation back to the planned schedule with minimized cost. Recovery options often include delaying flights, canceling flights, and changing (swapping) the aircraft for flights. The recovery horizon is generally one to four days.

Because of the vast decision space of the recovery problem and the quick response requirement, in real practice, airlines often decompose the recovery process into several stages and solve them in a sequential manner. Because aircraft are usually airlines' most expensive assets, the aircraft recovery problem is typically solved first. The ARP is to determine for each aircraft when and which flights to operate, with the objective of minimizing the total costs of flight cancellation, flight delay, and aircraft swap, while satisfying constraints such as maintenance, time and space matches, and airport capacity. Given the determined aircraft routes, in the next stage, a crew recovery problem is solved to re-dispatch crews to aircraft. Finally, passengers are re-accommodated by solving a passenger recovery problem. It should be noticed that the aircraft recovery problem is the fundamental stage of the whole recovery process: by and large, if fewer flights are cancelled and delayed in the first stage, fewer crew and passengers need to be re-dispatched and re-accommodated in the second and third stages, respectively.

There is a rich literature for the airline recovery problem (Clausen et al., 2010). Teodorović and Guberinić (1984) presented one of the pioneering studies in the airline recovery problem. They proposed a heuristic which solves each aircraft's route sequentially as a network flow problem using branch-and-bound. Jarrah et al. (1993) developed two network flow models, one for delay and one for cancellation. The models repeatedly solve shortest path problems for necessary flows. Argüello et al. (1997) proposed a greedy randomized adaptive search procedure (GRASP) to reconstruct aircraft routes. The algorithm follows a local search paradigm: first an incumbent solution is randomly selected from a candidate list and then neighbors of the incumbent solution are constructed; finally, the most desirable neighbor is put into the candidate list. The local search procedure repeats until a stopping criterion is met. Based on a time-space network of airports, flight legs and ground arcs, Yan and Lin (1997) devised network flow models that are solved by a network simplex method and a Lagrangian relaxation-based algorithm. Cao and Kanafani (1997a,b) modeled a quadratic zero-one programming for a multi-fleet recovery problem which is solved by an approximation linear programming (LP) algorithm. Similarly, Thengvall et al. (2001) considered a multi-fleet aircraft recovery problem after hub closures. Three models based on multi-commodity networks were presented. The authors further improved their work in Thengvall et al. (2003), where a bundle algorithm is applied to solving the model.

Different from network flow based models, Rosenberger et al. (2003) formulated a set partitioning model for rerouting aircraft with a heuristic for pre-selecting the aircraft which are to be rerouted. Selected aircraft are allowed to swap with disrupted aircraft and are included in a route generation procedure. Andersson and Värbrand (2004) developed an approach based on a set packing formulation which is derived from Dantzig-Wolfe decomposition. LP relaxation and a Lagrangian heuristic were proposed for the master problem; two column generation heuristics are implemented for the subproblems. However, maintenance is not considered in their model. Eggenberg et al. (2010) introduced constraint-specific recovery network for solving the problem. In the network, continuous timeline is discretized into time windows whose width is a parameter that needs to be tuned.

As mentioned above, after solving the aircraft recovery problem, airlines solve the crew and passenger recovery problems to obtain a complete recovery solution. In recent years, with the improvement of modeling approaches and computing capabilities, various methods have been developed to solve the two or three recovery problems in an integrated way. Nevertheless, still, the aircraft recovery problem plays a crucial part in the integrated methodologies. Petersen et al. (2012) published a fully integrated framework consisting of the three recovery problems using Benders decomposition, which is closely related to the work of Lettovsky (1997). To limit the size and complexity of the fully integrated problem, only pre-selected flights are input into the model for computation. With respect to the literature solving two recovery problems simultaneously, Abdelghany et al. (2008) utilized a simulation model and resource assignment optimization to solve the aircraft and crew recovery problems. Maher (2016) integrated the crew and aircraft recovery problems by column-and-row generation. Moreover, regarding the joint aircraft and passenger recovery problem, Bratu and Barnhart (2006) proposed two passenger recovery models in which passenger itinerary delays and cancellations are estimated in the formulation. A large neighborhood search heuristic was devised by Bisaillon et al. (2011). This heuristic is composed of three phases: construction, repair, and improvement, which iteratively destroys and repairs parts of the solution. Another heuristic-based framework for the joint aircraft and passenger recovery problem is by Jozefowiez et al. (2013). This heuristic also contains several stages; in the first stage, aircraft recovery is performed. Most recently, Zhang et al. (2016) developed a math-heuristic algorithm for recovering aircraft and passengers together. The algorithm carries out an aircraft recovery first; then, flights are re-scheduled and passengers are re-accommodated iteratively until a tolerance limit is reached.