European Journal of Operational Research 233 (2014) 234–245

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Innovative Applications of O.R.

Improvements to a large neighborhood search heuristic for an integrated aircraft and passenger recovery problem

Karine Sinclair*, Jean-François Cordeau, Gilbert Laporte

HEC Montréal and CIRRELT, 3000 chemin de la Côte-Sainte-Catherine, Montréal H3T 2A7, Canada

ARTICLE INFO

Article history:
Received 13 August 2012
Accepted 23 August 2013
Available online 13 September 2013

Keywords:
Airline recovery
Fleet assignment
Aircraft routing
Passenger itineraries
Large neighborhood search

ABSTRACT

Because most commercial passenger airlines operate on a hub-and-spoke network, small disturbances can cause major disruptions in their planned schedules and have a significant impact on their operational costs and performance. When a disturbance occurs, the airline often applies a recovery policy in order to quickly resume normal operations. We present in this paper a large neighborhood search heuristic to solve an integrated aircraft and passenger recovery problem. The problem consists of creating new aircraft routes and passenger itineraries to produce a feasible schedule during the recovery period. The method is based on an existing heuristic, developed in the context of the 2009 ROADEF Challenge, which alternates between three phases: construction, repair and improvement. We introduce a number of refinements in each phase so as to perform a more thorough search of the solution space. The resulting heuristic performs very well on the instances introduced for the challenge, obtaining the best known solution for 17 out of 22 instances within five minutes of computing time and 21 out of 22 instances within 10 minutes of computing time.

1. Introduction

In order to successfully manage their expensive resources, commercial passenger airlines must make use of efficient planning systems. For this reason, operational research plays a major role in the airline industry's tactical planning. As was documented by Barnhart, Belobaba, and Odoni (2003), this industry provides a favorable environment for the application of operations research (OR) models and techniques. The intensive use of computers also explains the importance of OR in the airline industry.

Because of their size and complexity, tactical planning problems are usually solved in a sequential order and are divided into five phases. Airlines must first determine which cities they will service and create a flight schedule. They then need to determine which type of aircraft will be assigned to the different flight legs of the schedule. In the third phase, airlines must create, for each aircraft, rotations that connect the flight legs while respecting maintenance constraints. Next, airlines must create crew pairings while respecting complex government and work related constraints. Finally, these crew pairings are combined to form monthly schedules for all crew members.

Since most commercial passenger airlines operate on a hub-and-spoke network, small disturbances can cause major disruptions in their planned schedules. Although disruptions can be caused by many factors such as cancelled or delayed flights, unavailable aircraft or crews, security measures and airport congestion, unfavorable weather conditions are the primary cause and according to Rosenberger, Johnson, and Nemhauser (2003), they are responsible for 75% of all disturbances. Disruptions have a significant impact on the operational costs and profits of airlines. Therefore, when they occur, it is necessary for the airlines to reestablish the planned schedule as quickly as possible, usually by the following day. Also, recovery problems need to be solved in a very short period of time, usually within minutes. Because of the size of the recovery problems and the time constraint, similar to the tactical planning, the recovery process is usually performed in a sequential way. First, the aircraft recovery problem is solved by either cancelling or delaying flights, modifying aircraft rotations or reassigning available aircraft while respecting maintenance, flow and location constraints. The second phase is crew recovery, which can be done by reassigning some crews, by deadheading crew members or by using reserve crews. Finally, the passenger recovery problem is solved. Although this sequential process can produce optimal solutions for the different phases, it rarely yields optimal solutions for the entire system. Hence, solving the problem globally, i.e. integrating the different recovery phases, should yield better solutions. Because disturbances occur frequently in the airline industry and their financial impact is substantial, it is necessary for the airlines to develop integrated approaches for the recovery problem.

---

* Corresponding author. Tel.: +1 5145250449.
E-mail addresses: karine.sinclair@hec.ca (K. Sinclair), jean-francois.cordeau@hec.ca (J.-F. Cordeau), gilbert.laporte@hec.ca (G. Laporte).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.08.034

Different methods have been developed to solve the aircraft recovery problem. Teodorović and Guberinić (1984) proposed a heuristic based on a network flow model and used a branch-and-bound method, while Jarrah and Yu (1993) also used network flow models, but their algorithm solves a shortest path problem repeatedly. Arguello, Bard, and Yu (1997) applied a greedy randomized adaptive search procedure (GRASP) consisting of a construction phase and a local search phase. Cao and Kanafani (1997a, 1997b) developed a linear programming approximation algorithm based on a quadratic programming model. Rosenberger et al. (2003) solved a set partitioning problem and used a heuristic to reduce the number of feasible routes. More recently, Eggenberg, Salani, and Bierlaire (2010) presented a constraint specific network recovery model which they solved by column generation. As for the crew recovery problem, Stojkovic, Soumis, and Desrosiers (1998) modeled the problem as a set partitioning problem while Lettovsky, Johnson, and Nemhauser (2000) and Medard and Sawhney (2007) modeled it as a set covering problem. Abdelghany, Ekollu, Narasimhan, and Abdelghany (2004) developed a mixed integer programming model. Other methods have been proposed to solve the crew recovery problem (see e.g. Nissen & Haase (2006) and Yu, Arguello, Song, McCowan, & White (2003)). To solve the passenger recovery problem, Bratu and Barnhart (2006) used network flow techniques, while Zhang and Hansen (2008) developed an integer non-linear programming model.

Although, to our knowledge, only Petersen, Solveling, Johnson, Clarke, and Shebalov (2010) addressed the full recovery problem (i.e. aircraft, crew and passenger recovery) and solved it using a Benders decomposition scheme, several methods have been developed to solve two recovery problems jointly. Abdelghany, Abdelghany, and Ekollu (2008) used a simulation model and a resource assignment optimization model to solve the joint aircraft and crew recovery problem. Luo and Yu (1997) modeled this joint problem as an integer linear program and solved it with a heuristic based on a restricted version of the model, while Stojkovic, Soumis, Desrosiers, and Solomon (2002) also modeled the problem as an integer linear program but solved it by network flow techniques. As for the joint aircraft and passenger recovery problem, Bratu and Barnhart (2006) presented two models and developed an Airline Operations Control Center simulator to evaluate their impact on a major US airline's operations. Their models also consider crew recovery, but only use an approximation of reserve crews and do not consider the disrupted crews. Zegordi and Jafari (2010) solved the aircraft recovery problem using an ant colony algorithm while taking into consideration disrupted passengers in the objective function. To solve the joint aircraft and passenger recovery problem, Jafari and Zegordi (2010) developed a mixed integer programming model. Finally, Ball, Barnhart, Nemhauser, and Odoni (2007) and Clausen, Larsen, Larsen, and Rezanova (2010) offer surveys devoted to airline recovery problems.

This paper presents a Large Neighborhood Search (LNS) heuristic for the joint aircraft and passenger recovery problem as described by Palpant, Boudia, Robelin, Gabteni, and Laburthe (2009) in the context of the 2009 ROADEF Challenge. Our method improves the LNS developed by Bisaillon, Cordeau, Laporte, and Pasin (2010) in the context of this challenge in which they won first prize. Eight other teams qualified for the final and their approaches are described on the web site http://chalenge.roadef.org/2009. Among the competing teams, three other approaches provided best solutions for at least two of the final instances. Jozefowiez, Mancel, and Mora-Camino (2010) developed a three-phase heuristic. The first phase integrates the disruptions to the initial plan and returns a feasible solution by removing flight legs and all disrupted itineraries. In the second phase, the cancelled itineraries from the previous phase are reassigned, when possible, to the existing rotations. Phase three attempts to create new flight legs to accommodate the remaining disrupted passengers using available aircraft. Mansi, Hanafi, Wilbault, and Clautiaux (2012), the team that finished second in the competition, proposed a two-phase heuristic based on an oscillation strategy and mathematical programming. This heuristic first attempts to find a feasible solution close to the initial schedule by solving a relaxation of the problem. If no feasible solution is obtained because of maintenance constraints, a dynamic programming based algorithm is used to find suitable routes. The second phase, the strategic oscillation, alternates between and destructive heuristics to improve solutions. The constructive phase generates feasible aircraft routes and passenger itineraries and then assigns them simultaneously to aircraft and passengers. The destructive phase deletes routes and cancels the corresponding passenger itineraries and then assigns the cancelled passengers to existing flights. Finally, Peekstok and Kuipers (2009) developed a simulated annealing algorithm. Initially, the algorithm accepts infeasibility with respect to airports, aircraft and passengers; however, once airport and aircraft feasibility is achieved, it does not allow the solution to become infeasible again for airports and aircraft. Although this team ranked sixth in the competition, it obtained two of the best solutions for the 22 instances considered in the final.

Acuna-Agost (2010) developed a network pruning algorithm that can be combined with aircraft recovery solution methods. The algorithm reduces the number of decision variables and constraints by identifying incompatible or suboptimal network nodes for each commodity. The algorithm was combined with the three-phase heuristic developed by Jozefowiez et al. (2010) to solve the problems introduced in the 2009 ROADEF Challenge.

The contribution of this paper is to introduce several refinements to the LNS algorithm that lead to a much improved performance. The resulting heuristic can quickly provide very good solutions to the problem. We also show that it can be profitable to run the algorithm for a longer period of time in order to accommodate additional passengers.

The remainder of the article is organized as follows. Section 2 provides a description of the joint aircraft and passenger recovery problem. Section 3 briefly describes the LNS heuristic developed by Bisaillon et al. (2010) and presents our improvements to this heuristic. Computational results are reported in Section 4, followed by the conclusion and discussion of future research directions in Section 5.

2. Problem description

The joint passenger and aircraft recovery problem consists of creating new aircraft routes and passenger itineraries so as to generate a feasible schedule during the recovery period and return to normal operations as quickly as possible. The problem can be represented on a time–space network $G = (N, A)$, where each node in the set $N = \{1, \ldots, n\}$ represents an airport at a specific time, and each arc in the set $A = \{(i, j);\ i, j \in N,\ i \neq j\}$ represents a flight leg or a connection between two flight legs at the same airport. Airports have restrictions on the maximum number of arrivals and departures allowed during each 60-minute period beginning on the hour. The parameter a_{ip} represents the arrival capacity at airport i during period p, whereas b_{ip} represents the departure capacity at airport i during period p.

The set F represents the fleet of aircraft operated by the airline, and each aircraft $f \in F$ is characterized by an identification number, a model and a cabin configuration. All aircraft of a given model have the same turn-round time, transit time, range