Les Cahiers du GERAD ISSN: 0711–2440

Simultaneous optimization of personalized integrated recovery for pilots and copilots

A. Kasirzadeh, F. Soumis,
F. Lessard, M. Saddoune

G–2015–65

July 2015

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Simultaneous optimization of personalized integrated recovery for pilots and copilots

Atoosa Kasirzadeh [a]
François Soumis [a]
François Lessard [a]
Mohammed Saddoune [a,b]

[a] *GERAD & Department of Mathematics and Industrial Engineering, Polytechnique Montréal, Montréal (Québec) Canada, H3C 3A7*

[b] *Université Hassan II, Faculté des Sciences et Techniques de Mohammedia, Mohammedia, Maroc*

atoosa.kasirzadeh@polymtl.ca
francois.soumis@gerad.ca
francois.lessard@gerad.ca
mohammed.saddoune@polymtl.ca

July 2015

Les Cahiers du GERAD
G–2015–65

Abstract: Various disturbances such as adverse weather conditions may result in delayed or canceled flights and affect the optimized schedules planned for airline crew members. In this paper, we solve the recovery problem via an integrated approach to reoptimize both the pairings and the personalized monthly plans. We solve this problem simultaneously for the pilots and copilots to obtain robust schedules that have the same pairings for pilots and copilots when possible. We propose a set partitioning formulation and we use column generation. We present results for seven instances from a major US carrier.

Key Words: Airline optimization, airline crew recovery, airline crew reoptimization, airline crew recovery, cockpit recovery.

Acknowledgments: This research has been supported by the Natural Sciences and Engineering Research Council of Canada and a collaborative R&D grant from AD OPT, a division of Kronos.

1 Introduction

Because of its complexity, the airline decision-making procedure is usually divided into *planning* and *recovery* stages (Klabjan, 2005; Belobaba et al., 2009). The planning stage is frequently further subdivided into *flight scheduling*, *fleet assignment*, *aircraft maintenance and routing*, and *crew scheduling* (Kasirzadeh et al., 2015). Crew scheduling is then separated into *crew pairing* and *crew assignment* (Barnhart et al., 2003; Gopalakrishnan and Johnson, 2005; Kasirzadeh et al., 2015). The crew pairing problem builds a minimum-cost set of pairings based on the scheduled flights such that the collective agreements and rules are respected. A *pairing* is a sequence of duties and overnight stops that starts and ends at a crew base. A *duty* is a sequence of flights (and/or deadheads) that forms a working day for a crew member; the duties are separated by overnight stops. Each crew member is associated with a *base* located at a large airport. A *monthly schedule* is a sequence of pairings separated by time off. The crew assignment problem combines the pairings, vacations, preassigned activities, and rest periods to build a set of monthly schedules that respect the regulations and the collective agreement. The assignment procedure is either *bidline* or *personalized*. In the bidline approach, anonymous monthly schedules are constructed and assigned to crew members. Personalized assignment is either a *rostering* or *seniority-based* procedure. The rostering approach aims to maximize the global satisfaction, whereas the seniority-based approach maximizes the satisfaction of the crew members in seniority order. Traditionally, the crew pairing and crew assignment problems have been solved sequentially; more recently, some researchers have integrated the two steps.

On the day of operation, external and/or internal perturbations may occur. These perturbations include late or absent crew members, aircraft breakdowns and unscheduled maintenance, security delays, air traffic control adjustments for meteorological reasons, and severe weather patterns such as snow storms. These disruptions result in delayed or canceled flights. Data from the Bureau of Transportation Statistics show that from 2005 to 2013 on average 20.39% of scheduled flights were delayed and 1.91% were canceled. Delayed and canceled flights directly affect the crew schedules, and adjustments become necessary. The recovery procedure is complex because it includes fleet reassignment and maintenance recovery, crew pairing and monthly schedule recovery, and passenger-itinerary recovery. These steps are often solved sequentially: first the flights are rescheduled, then the aircraft are rerouted, the crew schedules are updated, and the itineraries are adjusted. This traditional approach is presented in Figure 1.

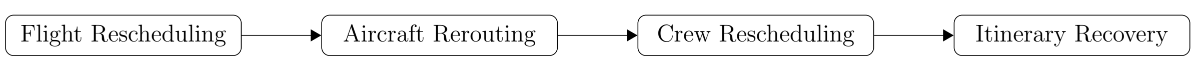

Figure 1: Schematic of sequential airline recovery procedure

In this paper, we focus on the crew rescheduling (recovery) problem, because the cost of the crew members is the second largest cost for airlines (after fuel). The algorithms for the crew recovery problem are similar to those applied for planning purposes. However, there are five major differences between the crew recovery and crew planning problems. First, the crew recovery problem cannot be separated into pairing and assignment steps. The updated pairings have to fit into the monthly schedules, so it is necessary to integrate the construction of new pairings and the adjustment of the monthly schedules. Second, the pilots and copilots must be treated simultaneously. The pairings should be the same for pilots and copilots, when possible, to maintain the robustness of the solution. When they are not the same, one flight perturbation will affect two different pairings, which will then affect more flights, and so the perturbation propagates through the monthly schedule. Third, the crew recovery problem must be solved quickly, whereas the crew planning problem is solved several weeks prior to operation. Fourth, the crew planning problem has a planning horizon that is frequently one month, whereas crew recovery reoptimizes the schedules locally for a period of a few days; therefore, the dimension of the optimization problem is reduced. Fifth, the objectives of crew planning are usually cost minimization and efficient crew utilization, whereas crew recovery has several conflicting objectives. These include minimizing the crew delays and minimizing the cost of the recovery operations.

The recovery problem must be small enough to be solved in a reasonable time, but its reoptimization domain must be sufficiently large to permit us to find feasible schedules for the rescheduled tasks. The main concern is to cover, in the most cost-efficient way, the set of flights while remaining as close as possible to the original schedules. It is important to minimize the number of flights that cannot be operated due to lack of sufficient crew. Crew recovery may involve rescheduling crew or deploying reserve crew members.

The contribution of this paper is an optimization approach for the integrated recovery of pairings and schedules for pilots and copilots simultaneously. This integrated approach considers both the pairing reoptimization and the recovery of monthly crew schedules, given all the relevant regulations. The problem is solved for pilots and copilots simultaneously to provide more robust schedules that reduce the propagation of perturbations. The rescheduled flights are input data. To the best of our knowledge, this paper presents the first mathematical programming method for the simultaneous recovery problem for pilots and copilots. Our main contribution is to demonstrate that a mathematical programming method can solve the personalized recovery problem for instances with up to 610 pilots and copilots in a reasonable time. We use column generation (CG).

The remainder of this paper is organized as follows. In Section 2, we provide a comprehensive literature review of crew recovery. Section 3 provides a detailed description of the problem, and Section 4 gives the mathematical formulation. Section 5 explains our algorithm, Section 6 gives our results, and Section 7 provides concluding remarks.

2 Literature review

Short computational times are required for airline recovery optimization, so the optimization problem must be small. We can either consider fewer crew members or restrict the time span of the reoptimization window.

To the best of our knowledge, the first survey of irregular airline operations is that by Clarke (1998b). He gives an extensive overview of the operations control center with respect to irregularities. He presents decision-support systems and algorithms based on operational data from the US domestic market. Filar et al. (2001) and Kohl et al. (2007) survey the state-of-the-art of decision-making for the airline recovery problem. They report on their research and development for large-scale airline disruption management. Another survey of the management of disruption in the airline industry is provided by Clausen et al. (2010). They also report a comparative study of aircraft/crew planning and recovery to explore the similarities between the solution approaches. Barnhart and Smith (2012) provide an overview of the role of OR in improving airline efficiency at the operational level.

Research on this topic began with investigations of aircraft scheduling in the presence of irregular operations (Clausen et al., 2010). This is a less complex problem: there are fewer aircraft than crew members and the aircraft rules are simpler than the crew-scheduling regulations. Teodorović and Guberinić (1984), Teodorović and Stojković (1990), Jarrah et al. (1993), Rakshit et al. (1996), Mathaisel (1996), Talluri (1996), Yan and Yang (1996), Clarke (1998a), Clarke (1997), Yan and Tu (1997), Cao and Kanafani (1997a), Cao and Kanafani (1997b), Luo and Yu (1997), Argüello et al. (1997), Luo and Yu (1998), Thengvall et al. (2000), Thengvall et al. (2001), Thengvall et al. (2003), Bard et al. (2001), Rosenberger et al. (2003), Andersson and Värbrand (2004), Andersson (2006), Liu et al. (2006), Liu et al. (2008), Eggenberg et al. (2007), and Zhao and Zhu (2007) studied aircraft recovery. We do not review this literature because the problem is not the focus of this paper.

There are three versions of the crew recovery problem. The first assumes that the flight schedules have already been recovered, i.e., the recovered flight schedules are input data for the crew recovery problem. Wei et al. (1997) and Song et al. (1998) provide a generalized set covering formulation for the crew pairing repair problem with reserve crew members. The objective is to repair the disturbed pairings as soon as possible while minimizing the cost. The branch and bound heuristic gives good results for small instances. Stojković et al. (1998) propose a set partitioning formulation for the operational crew scheduling problem and apply CG. They minimize the cost of covering all the flights with available crew members and minimize the crew disturbances. They allow only one modified pairing per crew member. To find the new pairings, they solve the crew pairing and the personalized monthly assignment problems simultaneously. They report results for small instances (with up to 32 crews and 210 flights) over one-day and seven-day periods. Medard and Sawhney (2007) expand the framework of Stojković et al. (1998) by permitting more than one modified pairing for each of the disrupted crew schedules. They propose an integrated pairing and assignment set covering problem in which the rescheduled flights replace the pairings. They solve the rescheduling problem by CG and provide results for small and medium instances. Nissen and Haase (2006) present a set covering formulation and a branch-and-price approach for the duty-period-based recovery problem for European airlines. They use CG and report results for small instances. Guo (2005) formulates the recovery problem as a set partitioning problem with the objective of minimizing the modifications to the planned schedule. CG and a genetic algorithm are used to find a balance between solution quality and computational time.

The second version of the problem allows flight cancellations. Johnson et al. (1994) present a set covering formulation. They take into account pairing and deadheading costs while forcing crew members to retain the same bases in the new solution. Results are presented for small instances. Lettovský et al. (2000) present a set covering formulation. They use a fast pairing generator and a branch and price technique, successfully handling small to medium disruptions. The pairing generation is designed to minimize the modifications to the original schedule. Yu et al. (2003) discuss the implementation of a crew recovery decision support system at Continental Airlines; it is a refined version of the model of Wei et al. (1997). They report good results and short computational times for small and medium instances.

The third version of the problem allows flight departures to be delayed. Stojković and Soumis (2001) extend the work of Stojković et al. (1998), presenting a set covering model and a CG approach. Reserve crew members are also allowed. They present results for instances with up to 59 pilots with 52 flights out of 190 being delayed. Stojković and Soumis (2005) extend this work. They simultaneously optimize the modifications to the flight departure times and the individual duties. The objective is to cover the maximum number of flights and to minimize the modifications to both the flights and the duties. Results for medium instances are reported. Abdelghany et al. (2004) provide a crew recovery decision support system for commercial hub-and-spoke airlines; they present good results for medium instances.

Since 1997 researchers have tried to integrate the different steps of the airline recovery problem. Lettovský (1997) presents an integrated approach for aircraft, crew, and passenger recovery and proposes an algorithm based on Benders decomposition. Bratu and Barnhart (2006) solve the passenger recovery problem while limiting the scheduling costs resulting from the perturbations. They permit delayed or canceled flights, and they make use of spare aircraft and reserve crew members. Zhang and Hansen (2008) present a model for a hub-and-spoke network that uses various modes of transportation to accommodate passengers whose travel plans have been perturbed. Abdelghany et al. (2008) present a commercial integrated approach to recovery when flights are delayed by severe weather conditions. Simultaneous recovery for aircraft and passengers is explored by Bisaillon et al. (2011) and Jafari and Zegordi (2011). Petersen et al. (2012) present a mathematical model and CG-based algorithm for the integrated flight, aircraft, crew, and passenger recovery problem. They give results for the hub-and-spoke network of a US carrier. Zhang and Lau (2014) present a set partitioning formulation for integrated flight, aircraft, and crew recovery. They provide a rolling-horizon algorithm and give results for small and medium instances from a US carrier.

3 Problem description

The goal of crew recovery is to quickly produce good solutions that cover the perturbed flights. Compared to the planning problem, crew recovery is more localized, focusing on the components that are affected by disturbances. Although only small portions of the crew schedules are affected, all the rules and regulations must continue to be satisfied for the full month. It is important to keep pilot-copilot pairs together in the new duties and pairings; this helps to ensure more robust schedules. Figure 2 shows the reoptimization window to illustrate the reduced size of the crew recovery problem.

Our primary goals for the simultaneous optimization of the pilot and copilot recovery problems are: (1) recovering the pairings and monthly schedules together, (2) recovering the pairings and monthly schedules for pilots and copilots simultaneously, and (3) solving the recovery problem quickly. We propose a heuristic that iterates between the *pilot recovery problem (PRP)* and *copilot recovery problem (CRP)*. At each iteration, the PRP or the CRP is solved using CG. The objective is to cover all the flights (perturbed and unperturbed) that lie within the reoptimization window while satisfying the preferences of the pilots and copilots, if possible.

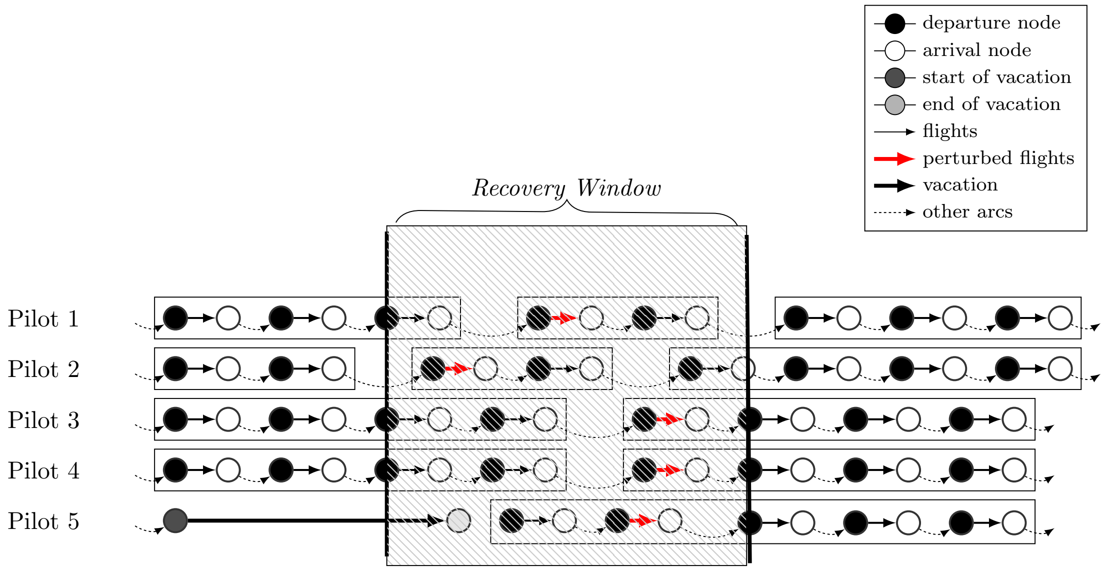

Figure 2: Reduced size of crew recovery problem

The algorithm starts from a set of monthly schedules for the pilots. In the first step, it takes the perturbed flights into account and solves the PRP over the reoptimization window. The set of pairings that lies within the recovery window is updated accordingly. This set of reoptimized pairings will fit within the monthly schedules for the pilots and copilots. In the second step, given this new set of pairings and the initial monthly schedules for the copilots, we solve the CRP. Using the new pairings obtained, we solve the PRP again and so on. This process continues until a stopping criterion is satisfied. We use a stopping criterion (a maximum number of iterations) because it may take a long time for the algorithm to converge. The algorithm is illustrated in Figure 3.

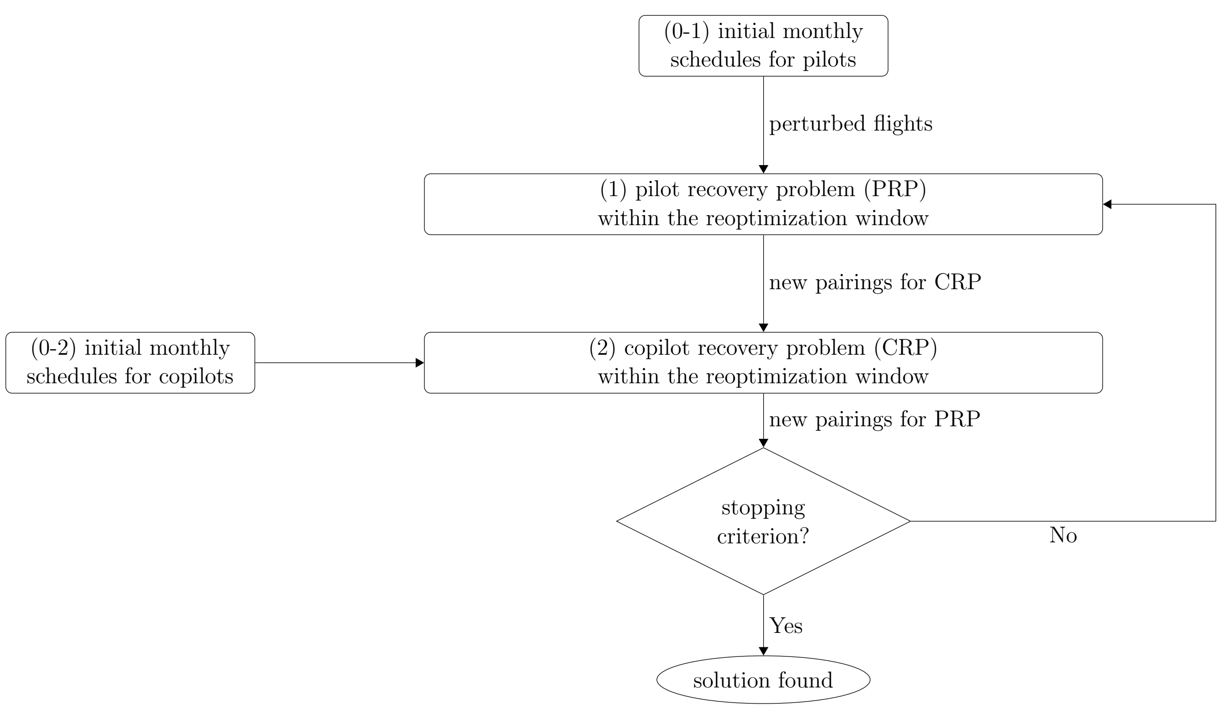

Figure 3: Heuristic algorithm for integrated crew recovery problem

4 Mathematical formulation

The simultaneous cockpit crew recovery problem is mathematically formulated using the following notation:

Sets

F_n: set of unperturbed flights in recovery window;
F_r: set of rescheduled flights;
P_r: set of feasible pairings overlapping recovery window;
L: set of pilots;
$V_{l,r}$: set of vacation preferences for pilot $l \in L$ in recovery window;
$G_{l,r}$: set of preferred flights for pilot $l \in L$ in recovery window;
S_l: set of feasible schedules for pilot $l \in L$;

Parameters

$C_{s,l}^{recovery}$: cost of personalized schedule $s \in S_l$ for pilot $l \in L$ during recovery window;
$\bar{C}_f$: penalty for not covering flight $f \in F_n \cup F_r$;
C_p: cost of feasible pairing $p \in P_r$;
$n_{s,r}^l$: number of preferred flights in recovery window in schedule $s \in S_l$ for pilot $l \in L$;
$c_{f,r}^l$: bonus for covering preferred flight $f \in G_{l,r}$ for pilot $l \in L$;
c_v^l: penalty for not covering vacation preference $v \in V_l$;

$$e_f^{s,l} = \begin{cases} 1 & \text{if flight } f \in F_n \cup F_r \text{ is covered by pilot } l \in L \text{ in schedule } s \in S_l \\ 0 & \text{otherwise;} \end{cases}$$

$$e_p^{s,l} = \begin{cases} 1 & \text{if pairing } p \in P_r \text{ is covered by pilot } l \in L \text{ in schedule } s \in S_l \\ 0 & \text{otherwise;} \end{cases}$$

$$v_v^{s,l} = \begin{cases} 1 & \text{if vacation } v \in V_{l,r} \text{ is covered by schedule } s \in S_l \\ 0 & \text{otherwise;} \end{cases}$$

Variables

$$x_l^s = \begin{cases} 1 & \text{if schedule } s \in S_l \text{ for pilot } l \in L \text{ is chosen} \\ 0 & \text{otherwise;} \end{cases}$$

$$\bar{e}_f = \begin{cases} 1 & \text{if flight } f \in F_n \cup F_r \text{ is not covered} \\ 0 & \text{otherwise.} \end{cases}$$

The recovery formulation for PRP is:

$$\min \sum_{l \in L} \sum_{s \in S_l} C_{s,l}^{recovery} x_l^s + \sum_{f \in F_n \cup F_r} \bar{e_f} \bar{C}_f \tag{1}$$

$$\text{s.t.} \sum_{l \in L} \sum_{s \in S_l} e_f^{s,l} x_l^s + \bar{e}_f = 1, \quad \forall f \in F_n \cup F_r \tag{2}$$

$$\sum_{s \in S_l} x_l^s \leq 1, \quad \forall l \in L \tag{3}$$

$$x_l^s \in \{0,1\}, \quad \forall l \in L, \forall s \in S_l \tag{4}$$

The recovery formulation for CRP is the same as (1)–(4) with the set of pilots (L) replaced by the set of copilots. The cost of the personalized schedule within the recovery window for pilot $l \in L$ is calculated via

$$C_{s,l}^{recovery} = \sum_{p \in P_r} e_p^{s,l} C_p + n_{s,r}^l c_{f,r}^l + \sum_{v \in V_{l,r}} (1 - v_v^{s,l}) c_v^l.$$

The objective (1) minimizes the total cost associated with the pilot schedules restricted to the recovery window. Constraints (2) ensure that perturbed and unperturbed flights within the reoptimization window are covered exactly once. Constraints (3) assign at most one schedule to each pilot, and constraints (4) are the binary requirements for the variables.

The schedule cost is composed of the pairing cost plus the penalties and bonuses for preferences. In practice, the pairing cost has a complex nonlinear structure, and an approximation is often used. To calculate the pairing cost, we use the cost function of Saddoune et al. (2013) that includes the deadhead, waiting, and duty costs. We take the preferences into account by introducing bonuses and penalties. Our preliminary results suggest a bonus of -50 for covering a preferred flight and a penalty of 5000 for not covering a vacation preference. The cost of not covering a scheduled flight is set to 10000 (all costs are in dollars). These values ensure that a good percentage of the preferences are satisfied while the gap remains small.

5 Algorithm

We use CG to solve the simultaneous personalized integrated recovery problem for pilots and copilots. CG is one of the most practical approaches for large-scale mixed integer problems (Klabjan, 2005). The linear relaxation of the recovery problem (1)–(4) is called the master problem. At each iteration of CG, we consider a restricted master problem (RMP) that contains a subset of the columns (variables). We solve the RMP using a standard linear programming algorithm such as the simplex method. This gives an optimal objective-function value and a pair of primal and dual solutions. Given this optimal dual solution, the current subproblem tries to find columns with negative reduced costs. If such columns are found, they are added to the RMP for the next iteration. Each subproblem corresponds to a resource-constrained shortest path problem, and we solve it by dynamic programming. When no variable with a negative reduced cost can be found, the optimal solution for the RMP is optimal for the master problem. In practice, because of slow convergence, the CG is often stopped before optimality is reached. Two parameters determine the CG stopping criterion. We stop the CG if in the last i iterations the objective value has decreased by less than a threshold α. These values are selected based on preliminary tests: i is set to 25 for instances 1–5 and to 10 for instances 6–7, and α is set to 0.001%. These choices greatly reduce the search domain for the optimization.

We can associate an acyclic network G, with node set N and arc set R, with each employee. To solve the subproblems, we must find shortest paths within these networks with negative reduced costs that satisfy the resource constraints. We use dynamic programming. We use the network structure, resources, and label-setting algorithm of Kasirzadeh et al. (2014).

We use two branching strategies at each node of the branch and bound tree. The first strategy fixes all the fractional values greater than a predetermined threshold to 1; we set the threshold to 0.85. The second strategy forces two flights to be consecutive in a pairing. We choose the branching strategy for a given node by computing a score for each strategy and choosing the strategy with the higher score.

6 Computational results

In this section, we present results for seven test instances. They are based on historical data for scheduled flights operated by short- and medium-haul aircraft in a major North American airline. The reoptimization is performed on monthly personalized schedules for pilots and copilots that are obtained by solving the personalized crew scheduling problem. For instances 1–3 (which are small), we find the monthly schedules by solving the simultaneous personalized integrated scheduling problems for pilots and copilots (Kasirzadeh et al., 2014). For the larger instances 4–7, we construct the personalized monthly schedules using the sequential approach presented by Kasirzadeh et al. (2015).

All the tests were executed on a Linux PC equipped with an Intel (R) Xeon (R) processor clocked at 2.93 GHz. The heuristic is coded in C++. We use the GENCOL column generation library (version 4.5) and the linear programming solver CPLEX 12.4.

We apply all the constraints used to construct personalized schedules for pilots and copilots at the planning level. Severe weather is the hypothetical disturbance that we consider. We construct four scenarios for disturbed flights. For each instance, we assume that the perturbations occur only in the largest base. The four scenarios are:

1. The perturbation occurs between 5 p.m. and 6 p.m. on the 15th of the month. It results in a delay of one hour for all the flights departing in this interval. The reoptimization window is from 4 p.m. on the 15th to 4 a.m. on the 16th.
2. This is identical to scenario 1 except that the perturbation results in a delay of two hours.
3. This perturbation affects 50% of the flights that arrive or depart between 4 p.m. and 6 p.m. on the 15th of the month. It results in a delay of two hours. The reoptimization window is from 1 p.m. on the 15th to 4 a.m. on the 16th.
4. This perturbation affects 50% of the flights that arrive or depart between 10 a.m. and 1 p.m. on the 15th of the month. It results in a delay of one hour. The reoptimization window is from 8 a.m. on the 15th to 4 a.m. on the 16th.

Tables 1–7 present the features of each reoptimization problem and the results for each perturbation scenario. The first five rows indicate the size of the reoptimization problem. The *no. of crew members* indicates the number of pilots and copilots with planned schedules. The *no. of active flights* is the number of flights within the recovery window. The *no. of active duties* and the *no. of active pairings* are the numbers of duties and pairings that overlap the reoptimization window and are included in the recovery problem. The *no. of delayed flights* is the number of delayed flights in the corresponding scenario.

We use three indicators to assess the algorithm: the *no. of CG iterations* indicates the total number of CG iterations for the three iterations of the heuristic. The *CPU time* (in seconds for instances 1–3 and minutes for instances 4–7) indicates the total CPU time for the three iterations. The *gap* is the percentage difference between the LP solution and the integer solution. All the instances are solved in a reasonable computational time. Except for instance 1 (the smallest instance), the gap is smaller than 1.03. In practice, for small instances that are not hard to solve, it is advisable to apply different branching strategies when the gap is greater than 1%.

We use six indicators to assess the solution quality. The *pairing similarity* is the percentage of common pairings for pilots and copilots within the recovery window at the final iteration of the reoptimization process. There are two ways to encourage common sets of duties and pairings. The first is to introduce *soft constraints*: penalties for duties and pairings that are different. The second is to introduce hard constraints; this restricts

Table 1: Results for Instance 1

	Scenario 1		Scenario 2		Scenario 3		Scenario 4	
	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots
No. of crew members	33	33	33	33	33	33	33	33
No. of active flights	22		22		30		34	
No. of active duties	12		12		13		13	
No. of active pairings	12		12		13		13	
No. of delayed flights	6		6		9		5	
Total no. of CG iterations	70	73	78	84	75	84	111	99
Total CPU time (s)	1.50	1.20	1.70	1.20	1.60	1.30	2.40	1.50
Gap (%)	0.00	2.24	0.00	1.91	0.00	1.19	0.00	0.00
Pairing similarity (%)	100		100		76.92		100	
Duty similarity (%)	100		100		76.92		100	
No. of uncovered flights	1	1	1	1	1	0	0	0
Loss of flight preferences (%)	-1.52	3.03	-1.52	-1.52	0.41	4.87	0.00	0.00
Loss of vacation preferences (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Cost increase (%)	0.98	0.98	1.36	1.37	-1.42	11.05	2.49	2.50

Table 2: Results for Instance 2

	Scenario 1		Scenario 2		Scenario 3		Scenario 4	
	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots
No. of crew members	34	34	34	34	34	34	34	34
No. of active flights	31		31		44		52	
No. of active duties	16		16		17		18	
No. of active pairings	15		15		16		17	
No. of delayed flights	5		5		4		7	
Total no. of CG iterations	66	63	66	63	140	115	129	147
Total CPU time (s)	0.90	0.90	0.90	0.90	3.00	2.20	3.00	4.20
Gap (%)	0.00	0.00	0.00	0.00	0.34	0.01	0.00	0.00
Pairing similarity (%)	80.00		80.00		100		100	
Duty similarity (%)	81.25		81.25		100		100	
No. of uncovered flights	0	0	0	0	0	0	0	0
Loss of flight preferences (%)	2.82	2.90	3.23	2.89	-0.59	0.33	0.08	0.04
Loss of vacation preferences (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Cost increase (%)	4.25	4.31	4.71	4.92	2.21	2.23	3.31	0.04

Table 3: Results for Instance 3

	Scenario 1		Scenario 2		Scenario 3		Scenario 4	
	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots
No. of crew members	47	47	47	47	47	47	47	47
No. of active flights	41		41		51		63	
No. of active duties	19		18		18		22	
No. of active pairings	22		21		20		24	
No. of delayed flights	6		6		4		10	
Total no. of CG iterations	114	88	87	72	122	205	121	126
Total CPU time (s)	7.00	7.20	6.60	5.40	8.50	17.01	10.5	13.80
Gap (%)	1.03	0.84	0.00	0.00	0.00	0.07	0.00	0.00
Pairing similarity (%)	95.45		95.24		85.00		100	
Duty similarity (%)	94.74		94.44		83.33		100	
No. of uncovered flights	1	0	0	1	0	0	0	0
Loss of flight preferences (%)	5.88	0.60	-0.08	1.96	4.48	3.24	3.04	5.54
Loss of vacation preferences (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Cost increase (%)	0.14	5.43	6.51	2.56	-1.88	-3.02	0.04	-0.04

Table 4: Results for Instance 4

	Scenario 1		Scenario 2		Scenario 3		Scenario 4	
	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots
No. of crew members	145	145	145	145	145	145	145	145
No. of active flights	125		125		162		188	
No. of active duties	55		57		55		66	
No. of active pairings	55		57		54		65	
No. of delayed flights	12		12		11		9	
Total no. of CG iterations	193	212	309	256	370	277	306	263
Total CPU time (min)	36.62	41.21	39.58	39.98	42.81	41.46	45.92	50.01
Gap (%)	0.00	0.00	0.09	0.10	0.13	0.02	0.00	0.00
Pairing similarity (%)	87.27		77.19		77.78		96.92	
Duty similarity (%)	90.91		80.70		78.18		96.97	
No. of uncovered flights	0	0	0	0	0	0	0	0
Loss of flight preferences (%)	-4.30	-1.48	-5.16	-2.18	-1.81	1.49	1.70	0.93
Loss of vacation preferences (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Cost increase (%)	5.43	2.21	4.91	1.89	2.38	0.27	1.02	1.14

Table 5: Results for Instance 5

	Scenario 1		Scenario 2		Scenario 3		Scenario 4	
	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots
No. of crew members	247	247	247	247	247	247	247	247
No. of active flights	115		115		149		182	
No. of active duties	88		88		96		103	
No. of active pairings	89		89		96		101	
No. of delayed flights	14		14		9		7	
Total no. of CG iterations	168	153	145	144	192	180	231	234
Total CPU time (min)	11.43	11.21	8.35	8.24	11.71	11.30	12.92	14.60
Gap (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00
Pairing similarity (%)	85.39		80.90		82.29		86.14	
Duty similarity (%)	96.59		87.50		96.88		94.17	
No. of uncovered flights	0	0	0	0	0	0	0	0
Loss of flight preferences (%)	0.00	4.99	0.00	5.07	0.00	-6.52	0.00	-8.64
Loss of vacation preferences (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Cost increase (%)	0.26	0.53	1.55	0.58	1.47	3.90	2.66	4.51

Table 6: Results for Instance 6

	Scenario 1		Scenario 2		Scenario 3		Scenario 4	
	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots
No. of crew members	223	223	223	223	223	223	223	223
No. of active flights	129		129		169		197	
No. of active duties	77		78		82		86	
No. of active pairings	77		78		80		84	
No. of delayed flights	13		13		12		8	
Total no. of CG iterations	181	192	258	240	212	231	262	270
Total CPU time (min)	6.25	7.99	9.11	9.84	6.74	8.50	9.45	11.43
Gap (%)	0.00	0.00	0.03	0.01	0.00	0.00	0.00	0.00
Pairing similarity (%)	92.21		92.31		95.00		95.24	
Duty similarity (%)	100		100		100		100	
No. of uncovered flights	0	0	0	0	0	0	0	0
Loss of flight preferences (%)	-0.51	-0.60	-0.10	-0.01	-3.20	-0.59	-1.60	-2.17
Loss of vacation preferences (%)	0	0	0	0	0	0	0	0
Cost increase (%)	4.95	4.92	5.87	5.84	2.86	2.85	2.14	1.15

Table 7: Results for Instance 7

	Scenario 1		Scenario 2		Scenario 3		Scenario 4	
	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots	Pilots	Copilots
No. of crew members	305	305	305	305	305	305	305	305
No. of active flights	160		160		212		253	
No. of active duties	113		113		121		127	
No. of active pairings	115		115		122		126	
No. of delayed flights	9		9		23		12	
Total no. of CG iterations	209	212	217	220	277	265	321	324
Total CPU time (min)	19.76	18.33	22.37	18.90	27.16	22.63	32.94	30.45
Gap (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Pairing similarity (%)	93.91		97.35		89.34		95.28	
Duty similarity (%)	97.35		97.39		92.56		92.06	
No. of uncovered flights	0	0	0	0	0	0	0	0
Loss of flight preferences (%)	-4.40	-3.37	-4.40	-3.37	5.53	4.18	5.14	6.02
Loss of vacation preferences (%)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Cost increase (%)	4.14	1.71	4.24	1.81	4.21	1.48	3.89	1.83

the domain of exploration. In this study, we use soft constraints. Our preliminary results show that an acceptable level of similarity is achieved when we set the penalty to 300.

The *duty similarity* is the percentage of common duties for pilots and copilots in the reoptimization window at the final iteration of the heuristic. The *no. of uncovered flights* indicates the number of flights in the reoptimization window that are uncovered despite the penalty imposed. The *loss of flight preferences* is the percentage of flight preferences lost after the perturbation, and the *loss of vacation preferences* is the percentage of vacation preferences lost. Negative losses indicate that more preferences are satisfied. The *cost increase* is the percentage increase in the cost of the portion of the monthly schedules within the recovery window.

For the test instances, the pairing similarity varies between 76.92% and 100% (with an average of 90.61% over all the tests). The duty similarity ranges between 76.92% and 100% (with an average of 93.33% over all the tests). This is an acceptable level of common pairings in a reoptimization context. In practice, the associated penalty could be lower or higher, depending on the importance that the airline attaches to common pairings and duties. Except in some of the scenarios for instances 1 and 3, all the flights are covered. The bonus for preferred flights gives an acceptable cost increase for the pilots and copilots and an acceptable loss of flight preferences after the reoptimization process. The bonus could be lower or higher depending on the importance that the airline attaches to flight preferences. We cannot yet provide an analysis of how changes to the bonus will impact the cost of the schedules; this is a complex situation. One direction for further research is a study of the relationship between the bonus and the costs of the schedules. In our tests, none of the vacation preferences are lost after the perturbation.

7 Summary and conclusions

We have proposed a new set partitioning formulation and a new heuristic that solves the integrated personalized recovery problem for pilots and copilots simultaneously. Our results indicate that the reoptimization algorithm covers the perturbed flights with an acceptable cost increase and loss of flight preferences. It can solve instances with up to 610 pilots and copilots in a reasonable computational time.

References

Abdelghany, A., Ekollu, G., Narasimhan, R., and Abdelghany, K. (2004). A proactive crew recovery decision support tool for commercial airlines during irregular operations. Annals of Operations Research, 127(1–4): 309–331.

Abdelghany, K.-F., Abdelghany, A.-F., and Ekollu, G. (2008). An integrated decision support tool for airlines schedule recovery during irregular operations. European Journal of Operational Research, 185(2): 825–848.

Andersson, T. (2006). Solving the flight perturbation problem with meta heuristics. Journal of Heuristics, 12(1–2): 37–53.

Andersson, T. and Värbrand, P. (2004). The flight perturbation problem. Transportation Planning and Technology, 27(2): 91–117.

Argüello, M., Bard, J., and Yu, G. (1997). A GRASP for aircraft routing in response to groundings and delays. Journal of Combinatorial Optimization, 1(3): 211–228.

Bard, J., Yu, G., and Argüello, M. (2001). Optimizing aircraft routings in response to groundings and delays. IIE Transactions, 33(10): 931–947.

Barnhart, C., Cohn, A.M., Johnson, E., Klabjan, D., Nemhauser, G., and Vance, P. (2003). Airline crew scheduling. In: Hall, R. (ed.). Handbook of Transportation Science, International Series in Operations Research & Management Science, vol. 56, Springer US, pp. 517–560.

Barnhart, C. and Smith, B. (2012). Quantitative Problem Solving Methods in the Airline Industry. Springer, New York

Belobaba, P., Odoni, A., and Barnhart, C. (2009). The Global Airline Industry. John Wiley & Sons.

Bisaillon, S., Cordeau, J.-F., Laporte, G., and Pasin, F. (2011). A large neighbourhood search heuristic for the aircraft and passenger recovery problem. 4OR, 9(2): 139–157.

Bratu, S. and Barnhart, C. (2006). Flight operations recovery: New approaches considering passenger recovery. Journal of Scheduling, 9(3): 279–298.

Cao, J. and Kanafani, A. (1997a). Real-time decision support for integration of airline flight cancellations and delays Part I: Mathematical formulation. Transportation Planning and Technology, 20(3): 183–199.

Cao, J. and Kanafani, A. (1997b). Real-time decision support for integration of airline flight cancellations and delays Part II: Algorithm and computational experiments. Transportation Planning and Technology, 20(3): 201–217.

Clarke, M.D.D. (1997). The airline schedule recovery problem. Dissertation, University of Montreal.

Clarke, M.D.D. (1998a). Development of heuristic procedures for flight rescheduling in the aftermath of irregular airline operations. Ph.D. thesis, Massachusetts Institute of Technology, Flight Transportation Laboratory.

Clarke, M.D.D. (1998b). Irregular airline operations: A review of the state-of-the-practice in airline operations control centers. Journal of Air Transport Management, 4(2): 67–76.

Clausen, J., Larsen, A., Larsen, J., and Rezanova, N.J. (2010). Disruption management in the airline industry: Concepts, models and methods. Computers & Operations Research, 37(5): 809–821.

Eggenberg, N., Bierlaire, M., and Salani, M. (2007). A column generation algorithm for disrupted airline schedules. Technical report, Ecole Polytechnique Federale de Lausanne.

Filar, J.A., Manyem, P., and White, K. (2001). How airlines and airports recover from schedule perturbations: A survey. Annals of Operations Research, 108(1–4): 315–333.

Gopalakrishnan, B. and Johnson, E.L. (2005) Airline crew scheduling: State-of-the-art. Annals of Operations Research, 140(1): 305–337.

Guo, Y. (2005). A decision support framework for the airline crew schedule disruption management with strategy mapping. In Operations Research Proceedings 2004, pp. 158–165. Springer.

Jafari, N. and Zegordi, S.H. (2011). Simultaneous recovery model for aircraft and passengers. Journal of the Franklin Institute, 348(7): 1638–1655.

Jarrah, A.I., Yu, G., Krishnamurthy, N., and Rakshit, A. (1993). A decision support framework for airline flight cancellations and delays. Transportation Science, 27(3): 266–280.

Johnson, E., Lettovský, L., Nemhauser, G.L., Pandit, R., and Querido, S. (1994). Final report to Northwest Airlines on the crew recovery problem. Technical report, The Logistic Institute, Georgia Institute of Technology, Atlanta, USA.

Kasirzadeh, A., Saddoune, M., and Soumis, F. (2014). Simultaneous optimization of personalized integrated scheduling for pilots and copilots. Technical Report G-2014-111, Cahiers du GERAD.

Kasirzadeh, A., Saddoune, M., and Soumis, F. (2015). Airline crew scheduling: models, algorithms, and data sets. EURO Journal on Transportation and Logistics, DOI 10.1007/s13676-015-0080-x.

Klabjan, D. (2005). Large-scale models in the airline industry. In: Desaulniers, G., Desrosiers, J., and Solomon, M. (eds.). Column Generation. Springer US, pp. 163–195.

Kohl, N., Larsen, A., Larsen, J., Ross, A., and Tiourine, S. (2007). Airline disruption management: Perspectives, experiences and outlook. Journal of Air Transport Management, 13(3): 149–162.

Lettovský, L. (1997). Airline operations recovery: An optimization approach. Ph.D. thesis, Georgia Institute of Technology.

Lettovský, L., Johnson, E.L., and Nemhauser, G.L. (2000). Airline crew recovery. Transportation Science, 34(4): 337–348.

Liu, T.-K., Jeng, C.-R., Liu, Y.-T. and Chang, Y.-H. (2008). Disruption management of an inequality-based multi-fleet airline schedule by a multi-objective genetic algorithm. Transportation Planning and Technology, 31(6): 613–639.

Liu, T.-K., Jeng, C.-R., Liu, Y.-T. and Tzeng, J.-Y. (2006). Applications of multi-objective evolutionary algorithm to airline disruption management. IEEE International Conference on Systems, Man and Cybernetics, 2006, 5: 4130–4135.

Luo, S. and Yu, G. (1997). On the airline schedule perturbation problem caused by the ground delay program. Transportation Science, 31(4): 298–311.

Luo, S. and Yu, G. (1998). Airline schedule perturbation problem: Landing and takeoff with nonsplitable resource for the ground delay program. International Series in Operations Research and Management Science, 9: 404–432.

Mathaisel, D.F.X. (1996). Decision support for airline system operations control and irregular operations. Computers & Operations Research, 23(11): 1083–1098.

Medard, C.P. and Sawhney, N. (2007). Airline crew scheduling from planning to operations. European Journal of Operational Research, 183(3): 1013–1027.

Nissen, R. and Haase, K. (2006). Duty-period-based network model for crew rescheduling in European airlines. Journal of Scheduling, 9(3): 255–278.

Petersen, J.D., Sölveling, G., Clarke, J.-P., Johnson, E.L., and Shebalov, S. (2012). An optimization approach to airline integrated recovery. Transportation Science, 46(4):482–500.

Rakshit, A., Krishnamurthy, N., and Yu, G. (1996). System operations advisor: A real-time decision support system for managing airline operations at United Airlines. Interfaces, 26(2):50–58.

Rosenberger, J., Johnson, E., and Nemhauser, G. (2003). Rerouting aircraft for airline recovery. Transportation Science, 37(4): 408–421.

Saddoune, M., Desaulniers, G., and Soumis, F. (2013). Aircrew pairings with possible repetitions of the same flight number. Computers & Operations Research, 40(3): 805–814.

Song, M., Wei, G., and Yu, G. (1998). A decision support framework for crew management during airline irregular operations. In Operations Research in the Airline Industry, pp. 259–286. Springer.

Stojković, M. and Soumis, F. (2001). An optimization model for the simultaneous operational flight and pilot scheduling problem. Management Science, 47(9): 1290–1305.

Stojković, M. and Soumis, F. (2005). The operational flight and multi-crew scheduling problem. Yugoslav Journal of Operations Research, 15(1): 25–48.

Stojković, M., Soumis, F., and Desrosiers, J. (1998). The operational airline crew scheduling problem. Transportation Science, 32(3): 232–245.

Talluri, K.T. (1996). Swapping applications in a daily airline fleet assignment. Transportation Science, 30(3): 237–248.

Teodorović, D. and Guberinić, S. (1984). Optimal dispatching strategy on an airline network after a schedule perturbation. European Journal of Operational Research, 15(2): 178–182.

Teodorović, D. and Stojković, G. (1990). Model for operational daily airline scheduling. Transportation Planning and Technology, 14(4): 273–285.

Thengvall, B., Bard, J., and Yu, G. (2000). Balancing user preferences for aircraft schedule recovery during irregular operations. IIE Transactions, 32(3): 181–193.

Thengvall, B., Bard, J., and Yu, G. (2003). A bundle algorithm approach for the aircraft schedule recovery problem during hub closures. Transportation Science, 37(4): 392–407.

Thengvall, B., Yu, G., and Bard, J. (2001). Multiple fleet aircraft schedule recovery following hub closures. Transportation Research Part A: Policy and Practice, 35(4): 289–308.

Wei, G., Yu, G., and Song, M. (1997). Optimization model and algorithm for crew management during airline irregular operations. Journal of Combinatorial Optimization, 1(3): 305–321.

Yan, S. and Tu, Y. (1997). Multifleet routing and multistop flight scheduling for schedule perturbation. European Journal of Operational Research, 103(1): 155-169.

Yan, S. and Yang, D. (1996). A decision support framework for handling schedule perturbation. Transportation Research Part B: Methodological, 30(6): 405–419.

Yu, G., Argüello, M., Song, G., McCowan, S.M., and White, A. (2003). A new era for crew recovery at Continental Airlines. Interfaces, 33(1): 5–22.

Zhang, D. and Lau, H. (2014). A rolling horizon based algorithm for solving integrated airline schedule recovery problem. Journal of Automation and Control Engineering, 2(4): 332–337.

Zhang, Y. and Hansen, M. (2008). Real-time intermodal substitution: Strategy for airline recovery from schedule perturbation and for mitigation of airport congestion. Transportation Research Record: Journal of the Transportation Research Board, 2052(1): 90–99.

Zhao, X. and Zhu, J. (2007). Grey programming for irregular flight scheduling. IEEE International Conference on Grey Systems and Intelligent Services, 2007, pp. 1607–1611.