

Iterative Airline Scheduling

Oliver Weide, David Ryan and Matthias Ehrgott

Department of Engineering Science

The University of Auckland

Private Bag 92019, Auckland, New Zealand

email: {o.weide, d.ryan, m.ehrgott@auckland.ac.nz}

January 17, 2007

Abstract

In airline scheduling a variety of planning and operational decision problems have to be solved. We consider the problems aircraft routing and crew pairing: Aircraft and crew must be allocated to flights in a schedule in a minimal cost way. Although these problems are not independent, they are usually formulated as independent mathematical optimisation models and solved sequentially. This approach might lead to a suboptimal allocation of aircraft and crew, since a solution of one of the problems may restrict the solution of the problem solved later. Also, when minimal cost solutions are used in operations, a short delay of one flight can cause very severe disruptions of the schedule later in the day. We generate solutions that incur small costs and are also robust to typical stochastic variability in airline operations. We solve the two original problems iteratively. Starting from a minimal cost solution, we produce a series of solutions which are increasingly robust. Using data from domestic airline schedules we evaluate the benefits of the approach as well as the trade-off between cost and robustness. We extend our approach considering the aircraft routing problem together with two crew pairing problems, one for technical crew and one for flight attendants.

Keywords: airline scheduling, crew pairing, aircraft routing, robust

1 Introduction

A sequence of planning problems must be solved in airline scheduling: First, marketing decisions in the *schedule design problem* determine the schedule of flights the airline operates. Each flight is specified by origin, destination, departure date, departure time and duration. Given the set of flights in a schedule, the solution of the *fleet assignment model* determines which flight is operated by which aircraft type. The objective is to maximise profit with respect to the number of available aircraft and other resource constraints. Next, the *aircraft routing problem* seeks a minimal cost assignment of available aircraft to the flights. A routing is assigned to each individual aircraft such that each flight is covered by exactly one routing. The routings must satisfy maintenance restrictions. Once aircraft types are assigned to flights, the aircraft routing problem can be solved for each aircraft type separately. In a similar way to the aircraft routing problem, the *crew pairing problem* (or *tour of duty problem*) allocates crew to flights in a minimal cost way. A set of generic crew pairings is constructed subject to many rules so that each flight is covered exactly once. Under the assumption that the crew is only allowed to operate a single aircraft type (which is usually the case for pilots for example) the crew pairing problem can also be solved separately for each aircraft type. The last of the planning problems is *crew rostering*. Based on the constructed crew pairings, a line of work is assigned to each individual crew member. The reader is referred to Klabjan [16] for a detailed description of the various airline scheduling problems.

In this paper two of the above problems, aircraft routing and crew pairing, are considered. The fleet assignment model is important for large airlines with multiple aircraft types. In the context relevant for this paper, the fleet can be regarded as homogeneous and fleet assignment can be omitted. The crew rostering problem can be viewed as a separate optimisation problem with no influence on the cost of the overall solution and is also not considered.

Traditionally, the aircraft routing problem is solved prior to the crew pairing problem, although the two problems are interdependent. One of the important causes of delays in airline operations is the availability of crew. This is particularly true if crew change from one aircraft to another during a duty period, especially if there is minimum ground time between the flights. If the incoming flight has been delayed for some reason, then both that aircraft and the aircraft to which the crew are changing will depart late. This propagation of delay can cause serious disruptions to the operation of the flight schedule. Ehrgott and Ryan [8] and Yen and Birge [28] have shown that the “robustness” of crew pairing solutions can be significantly improved if aircraft changes are made when ground time between the incoming and outgoing flights is much greater than the minimum ground time. This can be achieved in the crew pairing problem by penalising aircraft changes when ground time is short. Robust crew pairing solutions then have “crew following the same aircraft” as much as possible and changing aircraft only when ground time between flights is much longer than the minimum. In this sense, the robust crew pairing solution depends on the given aircraft routing solution. It is natural to ask if it is possible to improve the robustness of the crew pairing solution by permitting changes to the aircraft routing solution to encourage the “aircraft to follow the crew”. In this paper we investigate this question and show that it is indeed possible to reduce the cost of the crew pairing solution and simultaneously increase its robustness by considering the two problems together.

We formulate the aircraft routing problem and the crew pairing problem in one integrated model. This model yields one optimal solution for the two problems where the objective function is a weighted sum of cost and some value attached to robustness. However, the model is very hard to solve. Integration increases the complexity of both problems which are already \mathcal{NP} -hard individually (Garey and Johnson [11]). Decomposition methods are proposed in the literature to solve the integrated problem (see Mercier et al. [20]) but long computation times are needed to solve the model to optimality.

Instead of solving the integrated model, we propose to solve the two original problems iteratively. We start with a minimal cost crew pairing solution without taking aircraft routings into account. Then, in each iteration we solve the individual aircraft routing problem first, taking into account the current crew pairing solution. Then, given the aircraft routing solution we resolve the crew pairing problem. We only use the objective functions in both problems to pass information from the problem solved previously to generate more and more robust solutions. Hence the constraints are unaltered and the complexity of the two problems is not increased. We stop the process when the level of robustness can not be improved any further.

This procedure generates a series of feasible solutions for the integrated model with varying cost and robustness measure. The airline is not required to associate a cost with robustness but can study the trade-off between cost and robustness and then choose a solution they prefer to operate.

Applying this approach to various domestic airline schedules we obtain low cost solutions that are highly robust. The solutions obey all rules imposed by the airline and are ready to be implemented in practice. Various crew groups such as captains, first officers and flight attendants are required to operate an aircraft. We therefore extend the iterative approach and consider multiple crew groups at the same time. We generate solutions for the aircraft routing problem together with solutions for two different crew groups and compare these with solutions generated by applying the iterative approach to each crew group separately.

The paper is organised as follows: In Section 2 the problems are formally stated. We review the most recent approaches in the literature in Section 3. In Section 4 the iterative solution approach is presented and Section 5 concludes with numerical results.

2 Problem Formulation

In this section we define the crew pairing problem and the aircraft routing problem. We discuss how both problems can be solved efficiently and we formulate a model that integrates the two problems. The integrated model is capable of generating solutions that incur small costs and are also operationally robust. We start with the crew pairing problem.

2.1 The Crew Pairing Problem

Given a flight schedule, the *crew pairing problem* is defined as the problem of assigning crews to flights in the schedule such that each flight is operated by exactly one crew. A sequence of flights which can be flown by a crew on one work day is called a *duty period*. An alternating sequence of duty periods and rest periods is called a *crew pairing* (or *tour of duty*). Any crew pairing must start and end at the same crew base and is restricted by a number of rules such as rest time regulations or flying time restrictions. Costs are associated with each crew pairing. In the crew pairing problem we seek a minimal cost set of pairings that partition the flights in the schedule, i.e. each flight is contained in exactly one pairing.

The pairings can be represented as columns of a binary $m \times n^P$ matrix A^P where m is the number of flights in the schedule and n^P is the number of possible pairings. Entry $(a_{ij})^P, 1 \leq i \leq m, 1 \leq j \leq n^P$ equals 1 if flight i is contained in pairing j and 0 otherwise. With this matrix representation we formulate the crew pairing problem as a standard *set partitioning model*:

$$\begin{aligned} \text{Minimise} \quad & (\mathbf{c}^P)^T \mathbf{x}^P \\ \text{subject to} \quad & A^P \mathbf{x}^P = \mathbf{1} \\ & \mathbf{x}^P \in \{0, 1\}^{n^P}. \end{aligned} \tag{1}$$

The element c_j^P of $\mathbf{c}^P \in \mathbb{R}^{n^P}$ is the cost associated with pairing j . The decision variable $x_j^P \in \{0, 1\}$ is equal to 1 if pairing j is contained in the solution and 0 otherwise. The cost of a pairing is composed of flight time and duty time salaries, and meal, rest and travel allowances. *Base-constraints* are added to the standard model to consider *base strengths* at the crew bases. The base strength restricts the number of crew pairings that can start at a crew base in a particular week or on a particular day. With base-constraints included the model is referred to as *generalised set partitioning model* since these constraints usually have non-unit right hand sides.

Because the number of pairings n^P is too large to efficiently generate all of them, column generation and branch-and-bound techniques (i.e. branch-and-price) are identified in the literature as the most successful methods to solve the problem, which is \mathcal{NP} -hard. Using *column generation* only a small fraction of all possible pairings is considered initially and additional pairings (columns) are generated during the execution of the algorithm.

We model the schedule as a directed *flight network* where arcs represent flights and nodes represent departure or arrival of a flight. Additionally, a *connection-arc* is linking the arrival of one flight with the departure of another flight if the two flights can be operated consecutively by the same crew. In this network each crew pairing is represented by a path and the column generation problem can be solved by a resource constrained shortest path problem. The rules the pairing must obey are incorporated into the shortest path algorithm as resource constraints.

To solve (1), first the LP relaxation of (1) is solved using column generation. Fractions appearing in this solution are caused by different crew pairings competing for the same flights. To eliminate these fractions and obtain an integer solution for (1), a branch-and-bound algorithm with a *constraint branching* strategy is used (Ryan and Foster [23]). In this strategy two flights in the solution are chosen that are partially covered by different crew pairings. In one branch solutions are considered where both flights are contained in the same crew pairing. In the other branch both flights must not be operated by the same crew pairing. A special case of constraint branching is *follow-on branching* where two flights must be operated consecutively by the same crew in one branch and must not be operated consecutively by the same crew in the other branch.

2.2 The Aircraft Routing Problem

The *aircraft routing problem* (also referred to as *tail assignment*) is the problem of assigning aircraft to a given set of flights in a schedule. We seek one routing for each aircraft such that each flight of the schedule is contained in exactly one routing. Each routing is subject to maintenance requirements and other flying restrictions and the number of available aircraft is given. In the crew pairing problem on the other hand we seek an unknown number of generic pairings to partition all flights. Each particular aircraft must be assigned to one specific routing, similar to the *crew rostering problem* where a line of work is assigned to a particular crew member.

In a similar way to the crew pairing problem, routings can be represented as columns of a binary $(m + a) \times n^R$ matrix A^R where m is the number of flights, a the number of available aircraft and n^R the number of possible routings. The first m rows are defined similarly to those in A^P : The element $(a_{ij})^R$, $1 \leq i \leq m$, $1 \leq j \leq n^R$ equals 1 if flight i is contained in routing j and 0 otherwise. Additionally, the element $(a_{m+i,j})^R$, $1 \leq i \leq a$, $1 \leq j \leq n^R$ equals 1 if routing j is operated by aircraft i and 0 otherwise. The last a constraints are referred to as *generalised upper bound (GUB) constraints (or aircraft convexity constraints)* and ensure that each aircraft is assigned to exactly one routing. With this matrix representation the aircraft routing problem can be formulated in a similar manner to the crew pairing problem:

$$\begin{aligned} & \text{Minimise} && (\mathbf{c}^R)^T \mathbf{x}^R \\ & \text{subject to} && A^R \mathbf{x}^R = \mathbf{1} \\ & && \mathbf{x}^R \in \{0, 1\}^{n^R}. \end{aligned} \tag{2}$$

The element c_j^R of $\mathbf{c}^R \in \mathbb{R}^{n^R}$ is the cost associated with routing j . The decision variable $x_j^R \in \{0, 1\}$ is equal to 1 if routing j is in the solution and 0 otherwise. As costs of a routing *through-values* are usually considered. A through-value is revenue attached to a pair of flights if they are operated in sequence by the same aircraft. This accounts for passenger demand for direct connections (i.e. without the necessity of changing aircraft) on certain itineraries. In the context of this paper a set of connections that must be operated by the same aircraft is given by the airline and hence no through-values are considered. Since we also assume that the operational cost of all aircraft are identical, the aircraft routing problem reduces to a feasibility problem. We introduce penalty costs for the aircraft routings to influence characteristics of the aircraft routing solution, see Section 4.

This form of the set partitioning model is called the *rostering model* and is also \mathcal{NP} -hard. As in the crew pairing problem, the number of possible routings is very large and column generation techniques are used to obtain a solution for the problem. Using the same network formulation as the crew pairing problem, aircraft routings can also be represented as paths. The column generation problem is formulated as a resource constrained shortest path problem where we again only generate columns which satisfy all rules.

In contrast to the crew pairing problem we now branch on aircraft-flight pairs to obtain integer solutions. In one branch a particular aircraft is forced to operate a particular flight while in the other branch the aircraft is not allowed to operate this flight.

To justify this strategy we refer to the theory of *perfect matrices* which was first proposed by Padberg [21]. For a perfect matrix A and the problem $\min\{\mathbf{c}^T \mathbf{x} : A\mathbf{x} = \mathbf{1}, \mathbf{x} \geq 0, \mathbf{c} \in \mathbb{Z}^n\}$ there always exists an optimal integral solution vector \mathbf{x} . We first introduce some notation:

A graph G is called *complete* if every node is adjacent to every other node. The *chromatic number* of G is the minimal number of different colours needed to colour all nodes of G such that no adjacent nodes have the same colour. A *subgraph* G' of G is a subset of the nodes of G together with all arcs of G linking nodes in the subset. A *clique* is a complete subgraph. A graph G is called *perfect* if for every subgraph G' of G the chromatic number of G' is equal to the maximal cardinality of a clique in G' .

Let G_I denote the *intersection graph* associated with a matrix A . The nodes of G_I correspond to columns of A and two nodes are linked by an edge if the two corresponding columns have a common 1 in any row. The matrix A is called perfect if the associated intersection graph is perfect.

We investigate the submatrices of A^R that consist of the columns of A^R associated with a single aircraft. The intersection graphs of these submatrices are complete since all columns have

a common 1 in the aircraft convexity constraint. Every subgraph of a complete graph is also complete. Also, in any complete graph the chromatic number equals the cardinality of a maximal clique which is equal to the number of nodes in the graph. This results in the following theorem:

Theorem 1 *Each submatrix of A^R , that consists of the columns associated with a single aircraft, is perfect.*

Hence, fractions in the solutions of the LP relaxation of the aircraft routing problem are caused by different aircraft competing for the same flight.

2.3 Robust and Integrated Crew Pairing and Aircraft Routing Problem

If two flights can be operated in sequence by the same crew or aircraft (i.e. there exists a connection-arc linking both flights), the time between arrival of the first and departure of the second flight is called *turn-time* for aircraft and *sit-time* for crew.

The minimal time required for an aircraft or crew to operate a connection is called *minimal turn-time* or *minimal sit-time*, respectively. The minimal sit-time can exceed the minimal turn-time. If the crew stays on the same aircraft both are identical. A connection between flights i and j is called *short* if

$$(\text{minimal turn-time})_{ij} \leq (\text{sit-time})_{ij} \leq (\text{minimal sit-time})_{ij}.$$

Thus, in a feasible solution, short connections are only allowed if the crew stays on the same aircraft. This condition might result in suboptimal or infeasible solutions if the two problems are solved separately.

Also, we prefer solutions where crew are not changing aircraft when the turn time is less than some *restricted time*. A connection between two flights i and j is called *restricted* if

$$(\text{minimal sit-time})_{ij} \leq (\text{sit-time})_{ij} \leq (\text{restricted time})_{ij}.$$

In contrast to short connections, crews are allowed to change aircraft if the connection is restricted, but we try to find solutions in which this occurs as rarely as possible. If crew change aircraft on restricted connections we refer to these connections as *restricted aircraft changes*.

It is unlikely that the schedule is operated as planned. Delays occur frequently in airline operations and can for example be caused by late passengers, unscheduled maintenance requirements or bad weather. Minimal turn-times are usually used in aircraft routings to keep costs low and connection times attractive for passengers. Hence, if a flight is delayed, the flight operated next by the aircraft is probably also delayed. But if the crew are also changing aircraft on a restricted connection after the delayed flight, at least one more flight might be affected by the initial delay. Due to the small buffer to compensate for the delay, the crew are likely to be late for the next flight they operate. This behaviour can propagate to a large number of delayed flights in a short amount of time.

A solution, where effects of potential delays are minimal, is called *operationally robust*. The concept of robust solutions is important since an airline is interested in achieving high *on-time performance (OTP)*, i.e. the percentage of all flights in the schedule that depart on-time. Bad OTP can incur large additional costs (referred to as *recovery costs*), caused by additionally required crews, compensations for passengers affected by delayed or cancelled flights and damaged reputation of the airline. These costs may by far exceed the savings of using a solution with less *planned cost* than using a solution that is more robust. We try to identify solutions with low planned costs which are operationally robust, i.e. where disruptions will result in minimal recovery costs. Costs listed in this paper are generally planned costs. We refer to the sum of planned costs and recovery costs as *operational costs*.

In Figure 1 examples of a non-robust (top) and a robust solution (bottom) are depicted. Flights are represented as rectangles with origin and destination airports indicated by 3-letter-codes. Aircraft routings are represented as rows of flights while crew pairings are represented as flights connected by lines. Dashed rectangles represent flights that are delayed. We can see in the first

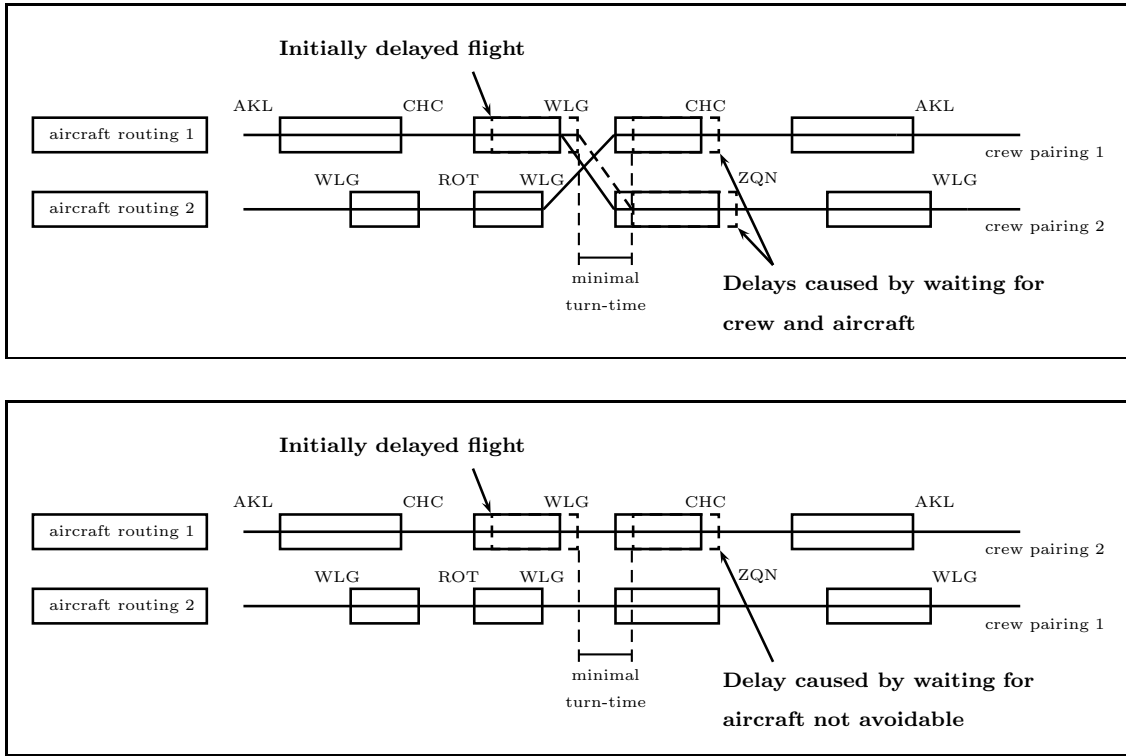


Figure 1: Comparison of a non-robust and a robust solution

scenario that two flights are effected by the initial delay because the crew operates a restricted aircraft change. In the second scenario only one other flight is effected by the initial delay.

In order to integrate the concepts of short and restricted connections into our formulation we enumerate all short and restricted connections. We define a binary $m^B \times n^P$ matrix B^P where m^B is the number of short connections. Each pairing is associated with one column of B^P , where $(b_{ij})^P, 1 \leq i \leq m^B, 1 \leq j \leq n^P$ equals 1 if short connection i is contained in pairing j and 0 otherwise. Analogously for restricted connections, we define a binary $m^D \times n^P$ matrix D^P where m^D is the number of restricted connections. $(d_{ij})^P, 1 \leq i \leq m^D, 1 \leq j \leq n^P$ equals 1 if restricted connection i is contained in pairing j and 0 otherwise. For aircraft a binary $m^B \times n^R$ matrix B^R and a binary $m^D \times n^R$ matrix D^R are defined in an analogous way.

With this matrix representation the *robust and integrated crew scheduling and aircraft routing problem* can be formulated as follows:

$$\begin{aligned}
 & \text{Minimise} && (\mathbf{c}^P)^T \mathbf{x}^P + (\mathbf{c}^R)^T \mathbf{x}^R + (\mathbf{c}^D)^T \mathbf{d} \\
 & \text{subject to} && A^P \mathbf{x}^P &= & \mathbf{1} \\
 & && && A^R \mathbf{x}^R &= & \mathbf{1} && (3) \\
 & && B^P \mathbf{x}^P - B^R \mathbf{x}^R &\leq & \mathbf{0} \\
 & && D^P \mathbf{x}^P - D^R \mathbf{x}^R - \mathbf{d} &\leq & \mathbf{0},
 \end{aligned}$$

where $\mathbf{x}^P \in \{0, 1\}^{n^P}$, $\mathbf{x}^R \in \{0, 1\}^{n^R}$, $\mathbf{d} \in \{0, 1\}^{m^D}$ are binary variables and $\mathbf{c}^D \in \mathbb{R}_+^{m^D}$ is some positive penalty parameter. Variable d_i equals 1 if restricted connection i is operated by a crew but no aircraft and 0 otherwise. The first two sets of constraints are identical to the original problem formulations. The third set of constraints enforces that short connections which are operated by some crew are also operated by some aircraft. The last set of constraints provokes additional cost

in the objective function if a restricted connection is operated by a crew and no aircraft. In Section 5 we find solutions for the integrated model applied to schedules of dated periods of one week.

3 Literature

Airline scheduling problems have been addressed in an extensive number of publications, see Klabjan [16] for a detailed overview. We list selected literature addressing airline scheduling problems and describe attempts to integrate the aircraft routing and crew pairing problems. We conclude with recent formulations that incorporate robustness measures.

We are not aware of contributions that discuss the schedule design problem as a single problem. Since schedule design is addressed in combination with fleet assignment we start with the literature on fleet assignment. The fleet assignment model is formulated by Hane et al. [15] as a multi-commodity flow problem. It is called *leg-based* because revenue effects between flight-legs are not modelled. To take such network effects into account, Barnhart et al. [1] describe an enhanced model using demand forecasts for origin-destination pairs.

An integrated model for schedule design and fleet assignment is presented by Lohatepanont and Barnhart [18]. They use the origin-destination fleet assignment model and flights are chosen from an optional set of flights to maximise profit.

The aircraft routing problem has been addressed in a number of publications, for example in Clarke et al. [4], Feo and Bard [10], Daskin and Panayotopoulos [7] and Gopalan and Talluri [13] and more recently in Bartholomew-Biggs et al. [3] and Grönkvist [14]. Sarac et al. [25] consider the aircraft routing problem on an operational level rather than a planning level. They add constraints to the set partitioning formulation to satisfy maintenance availability at the maintenance bases. Due to these constraints they must alter the branching strategy described in Section 2.2. They use a combination of follow-on (see Section 2.1) and aircraft-flight pair (see Section 2.2) branching.

After fuel costs, crew salary is the second largest operational cost an airline has to account for. Therefore finding a minimal cost solution to the crew pairing problem is important. On the other hand it is a very complicated problem due to the large number of possible pairings, the complicated rule structure and the necessity to find integer solutions. For these reasons the crew pairing problem has received a lot of attention in the literature, see Barnhart et al. [2] for a detailed description of the crew pairing problem and a review of the literature addressing the problem. Also recently, Gopalakrishnan and Johnson [12] give a comprehensive overview on state-of-the-art methods to solve the crew pairing problem.

The crew rostering problem can be viewed as a separate optimisation problem with no effect on the cost of the integrated solution. We refer to Ernst et al. [9] for an annotated bibliography of rostering problems.

Various combinations of airline scheduling problems have been integrated in the literature. Here we focus on the integration of crew pairing and aircraft routing problems.

Klabjan et al. [17] partially integrate aircraft routing, crew pairing and schedule design. They reverse the order of the crew pairing and aircraft routing problems. Plane count constraints are added to the crew pairing problem to ensure the existence of a feasible solution for the aircraft routing problem. Plane count constraints ensure that at most the number of available aircraft is used at any time. Their results are based on a *hub-and-spoke network*. In this network only large airports (hubs) are linked by direct flights and all smaller airports (spokes) are only connected to one hub. Many aircraft meet at the same hub at the same time ensuring the existence of many feasible connections. This property leads to a much larger number of feasible routings than in an *interconnected network*. In interconnected networks many airports are linked with multiple other airports by direct flights. To include schedule design in the model the departure time of each flight is allowed to vary in some time window. This is done by relaxing feasibility parameters in the crew pairing problem and hence generating a larger set of pairings. Each feasible pairing then has a departure time attached to each flight contained in the pairing. Klabjan et al. [17] solve the crew pairing problem via a linear program (LP) based branch-and-bound algorithm.

Another model to integrate aircraft routing and crew pairing is proposed by Cordeau et al. [6] and also by Mercier et al. [20]. They use Benders decomposition and branch-and-price to solve the model. Employing the crew pairing problem as the subproblem as well as the master problem has

been tested, the latter with better success. Both approaches add inequalities to the set partitioning polytopes of the problems. Cordeau et al. [6] also reverse the sequential approach and try to solve the crew pairing problem first followed by the aircraft routing problem (Klabjan et al. [17]). They apply this approach to an interconnected network but were not successful in obtaining feasible solutions for the aircraft routing problem.

Cohn and Barnhart [5] also integrate aircraft routing and crew pairing. They extend the crew pairing problem by using the aircraft routing problem as a second column generator next to the crew pairing generator. For each solution of the aircraft routing problem one variable is added to the crew pairing problem and a convexity constraint ensures the selection of one of the aircraft routing solutions in the final solution of the problem. LP based branch-and-price is used in this computationally expensive solution method. Mercier et al. [20] find that their Benders decomposition approach yields better solutions in less computation time than the extended crew pairing model of Cohn and Barnhart [5].

Sandhu and Klabjan [24] partially integrate fleet assignment, aircraft routing and crew pairing with a similar approach as Klabjan et al. [17] and solve the model with both Lagrangian relaxation and Benders decomposition.

Very recently, Mercier and Soumis [19] extend their model (Mercier et al. [20]) and integrate aircraft routing and crew pairing with time windows for the departure times. Flights are allowed to depart five minutes earlier or later than originally scheduled. Binary variables are used to indicate which departure time is assigned to a flight. Equality constraints sum up the binary departure time variables for the crew and aircraft solutions and ensure that the same departure times are used in the solutions of both problems. Again, the authors use Benders decomposition to solve the problem.

Models that focus purely on minimising cost tend to generate solutions that appear brittle in operations. Such solutions incur large recovery costs once disruptions occur in practice. In order to improve the behaviour in operations a number of robustness measures have been introduced.

Schaefer et al. [26] use expected operational cost for the crew pairings instead of planned cost. Interactive effects between pairings are ignored and a push-back strategy for recovery is used. In this strategy the flights are delayed until crew and aircraft are available. The authors use SimAir to estimate the costs and to evaluate the quality of their solutions. SimAir is a Monte Carlo simulation of airline operations that permits the evaluation of schedules and recovery policies in operations, see Rosenberger et al. [22].

Yen and Birge [28] formulate the crew pairing problem as a stochastic programming problem in a computationally expensive approach. Crew switching aircraft are penalised in the objective function. A similar measure of robustness is introduced by Ehrgott and Ryan [8] in a deterministic approach. Crew pairings are penalised where crew are changing aircraft and the sit-time of the crew is less than the minimal sit-time plus some measure of delay of the incoming flight. Crew who stay on the same aircraft are not penalised. Thus, crew connections where disruptions are likely to propagate onto multiple flights are penalised. Robustness is treated as a second objective function in a bicriteria approach.

Most recently, Shebalov and Klabjan [27] solve the crew pairing problem first and then maximise the number of move-up crews without increasing the planned cost too much. Move-up crews are crews that can potentially be swapped in case one crew is delayed. They compare their method with the method of solving the standard crew pairing problem by simulating disruptions and find solutions with significantly lower operational costs if the additional cost allowed for move-up crews is not too high.

4 Iterative Solution Approach

Currently, the most successful approach in the literature to solve the integrated crew pairing and aircraft routing problem seems to be Benders decomposition (see Mercier et al. [20]). However, large computation time is needed and the problem is only solved to near optimality.

We assume the aircraft routing cost to be fixed and consider the cost of the crew pairings as the only cost of the integrated solution. We consider a connection to be restricted, if the sit-time does not exceed the minimal sit-time by more than 30 minutes. The non-robustness measure (NRM) of

an integrated solution is a sum of sit-time dependent penalties over all restricted aircraft changes. We choose penalties which increase linearly with decreasing sit-time:

$$NRM = \sum_{ij \in RAC} (k_1 - ((\text{sit-time})_{ij} - (\text{minimal sit-time})_{ij})) * k_2, \quad (4)$$

where RAC is the set of all restricted aircraft changes. Weights k_1 and k_2 are chosen such that the weight equals 7 for a restricted aircraft change with sit-time equal to the minimal sit-time, 6 for a restricted aircraft change with sit-time exceeding the minimal sit-time by 5 minutes, and so on until a weight of 1 is assigned to a restricted aircraft change with sit-time exceeding the minimal sit-time by 30 minutes. Note that all departure and arrival times in the schedules we consider are discretised in 5 minute intervals. A different set of connections (e.g. those with a larger sit-time) or a different function of the sit-time (e.g. where weights increase exponentially with decreasing sit-time) could be chosen in a straightforward way.

We search for an integrated solution with small cost and small non-robustness measure. We do not attempt to solve the integrated problem to optimality. Instead we propose to solve the crew pairing problem and the aircraft routing problem iteratively. Initially we solve the crew pairing problem to cost optimality without considering any aircraft routings. Not considering any aircraft routings results in the generation of a larger set of feasible crew pairings since feasibility parameters are relaxed (see Section 2.3). This solution yields a lower bound on the cost of a feasible integrated solution. Then, in each iteration the aircraft routing problem is solved first. We consider all restricted connections operated in the current crew pairing solution and force the aircraft routing solution to contain as many of those connections as possible. This will enforce the ‘‘aircraft to follow the crew’’ as much as possible if the connection is restricted. In other words, we solve the aircraft routing problem using the following objective function:

$$\max \sum_{ij \in RC} ((k_1 - ((\text{sit-time})_{ij} - (\text{minimal sit-time})_{ij})) * k_2) x_{ij}, \quad (5)$$

where RC is the set of restricted connections operated in the current crew pairing problem and x_{ij} equals 1 if the aircraft routing solution contains connection ij and 0 otherwise and k_1, k_2 are chosen as above. Note that we maximise over the set of restricted connections. This is in contrast to the non-robustness measure (4) where we minimise over the set of restricted aircraft changes. Next we solve the crew pairing problem to optimality for the current aircraft routing solution with a weighted sum objective function of crew pairing costs and non-robustness measure:

$$\min \text{crewPairingCost} + \text{crewPenalty} \times NRM. \quad (6)$$

Hence each iteration yields a feasible solution to the integrated problem. We start with crewPenalty equal to zero and increase this weight in each iteration in order to increase the robustness of the solutions we generate.

The advantages of this procedure are as follows: Firstly, we can generate solutions with small cost and small non-robustness measure in a short amount of computation time and with no increase of complexity compared to the original problems. Secondly, we can study the trade-off between crew pairing costs and non-robustness measure of the solutions and the airline can choose a preferred solution without attaching a cost to the non-robustness measure a priori. Finally, we obtain a lower bound for the crew pairing cost which allows us to measure the quality of the final solution. Algorithm 1 shows the steps of the iterative approach.

For the schedule data sets relevant to this work, it can be assumed that the minimal sit-time of the crew is equal to the minimal turn-time of aircraft. Hence no short connections are taken into account. Since we always solve the crew pairing problem for a given solution to the aircraft routing problem, short connections could be considered by removing connections in the underlying network structure of the crew pairing problem. If short connections are present in the problem, Step 2 of Algorithm 1 generally yields an infeasible solution that violates the short connection rules (Section 2.3). The interdependence between aircraft routings and crew pairings stated in Section 2.3 is extended by an additional rule. We understand that this rule has been imposed by the airline in an attempt to improve robustness. The number of times crew are allowed to change aircraft in

Algorithm 1 Iterative Algorithm

- 1: SET $crewPenalty = 0$
 - 2: SOLVE crew pairing problem with objective function (6){Since no aircraft routings are taken into account a larger set of feasible pairings is generated.}
 - 3: **while** $crewPenalty \leq maxCrewPenalty$ **do**
 - 4: SOLVE aircraft routing problem with objective function (5){Maximise the number of restricted connections contained in the aircraft routing solution that are operated in the current crew pairing solution.}
 - 5: SOLVE crew pairing problem with objective function (6){Minimise cost and the number of restricted aircraft changes.}
 - 6: INCREASE $crewPenalty$
 - 7: BREAK if the non-robustness measure can not be improved
 - 8: **end while**
-

a single duty period is restricted to one for technical crew and two for flight attendants. In the following discussion we refer to these restrictions by DPACLIM (duty period aircraft change limit).

The cost of the crew pairing solution of Step 2 yields a lower bound on the crew pairing costs of a feasible integrated solution since no aircraft routings are taken into account. In our experiments, no aircraft routing solution for the crew pairing solution of Step 2 could be found to satisfy the DPACLIM rule. Hence, Step 2 usually yields an infeasible solution to the integrated problem. If the DPACLIM rule is relaxed we can find an aircraft routing solution to form a feasible solution to the integrated problem. For the test instances, such an integrated solution contains a large number of restricted aircraft changes and hence accounts for a large non-robustness measure.

After the initial steps of the algorithm, we obtain a feasible solution to the integrated problem in each iteration by solving the crew pairing problem (Step 5) for a given aircraft routing solution (Step 4). Once the integrated solution converged to a stable solution, the algorithm stops. For a stable solution, successive iterations yield identical aircraft routing and crew pairing solutions, despite increasing the $crewPenalty$ weight. Hence, the non-robustness measure can not be improved. The value of $maxCrewPenalty$ is chosen such that the non-robustness measure dominates the crew pairing costs in function (6) in the sense that no restricted aircraft changes are contained in the optimal solution if such a solution exists.

From a computational point of view, the main advantage of this solution algorithm to the integrated problem is that the original set partitioning formulations of both problems remain unchanged. This preserves the properties of the two problems and enables us to solve each problem efficiently with the methods described in Section 2.

If aircraft routing and crew pairing solvers are available, the implementation of the iterative approach is straightforward. We influence the characteristics of the solutions to either problem via the costs of the underlying network structure. This is easily implemented into the shortest path computations of the column generators for both problems: Penalties are applied to restricted connections used by the current crew solution in the network for the aircraft routing calculations. Since there are no other costs associated with the aircraft routings these penalties will enforce paths in the solution that contain as many restricted crew connections as possible. Also, we always obtain a feasible aircraft routing solution since all changes to the problem are in the objective function. In the crew pairing problem, each restricted aircraft change will incur a penalty during the shortest path calculation. In Section 5.4 we apply the iterative algorithm to a number of different data sets.

The crew for an aircraft usually consists of multiple crew groups. Our data sets consist of three crew groups, namely captains and first officers (i.e. technical crew) and flight attendants. Different rules, base strengths and pay structures apply to each group. A robust aircraft routing and crew pairing solution for one crew group may enforce many restricted aircraft changes in a crew pairing solution for another crew group. Hence, considering aircraft and one crew group might lead to a suboptimal solution. Ideally we want to consider all crew groups and the aircraft simultaneously. To incorporate multiple crew groups into the iterative approach we simply solve the crew pairing problem in Step 5 for each crew group separately but use the same common aircraft routing solution. To obtain penalties for the restricted connections for the subsequent aircraft routing problem we weight the penalties for the restricted connections of the different crew

pairing solutions according to weights (w_1 and w_2 in example 7) chosen by the airline. For example, the aircraft routing objective function in Step 4 for two crew groups is changed to:

$$\begin{aligned} \max \quad & w_1 \sum_{ij \in RC_1} ((k_1 - ((\text{sit-time})_{ij} - (\text{minimal sit-time})_{ij})) * k_2) x_{ij} + \\ & w_2 \sum_{ij \in RC_2} ((k_1 - ((\text{sit-time})_{ij} - (\text{minimal sit-time})_{ij})) * k_2) x_{ij}, \end{aligned} \tag{7}$$

where RC_1 and RC_2 are the sets of restricted connections of the crew pairing solutions for crew group 1 and crew group 2, respectively. Weights w_1 and w_2 can be chosen to reflect the ratio of crew pairing costs between both crew groups. We then generate a new aircraft routing solution as before. The results of the *iterative algorithm with multiple crew groups* are presented in Section 5.5.

5 Computational Experiments

5.1 Flight schedule

The computational experiments are performed on interconnected flight networks corresponding to domestic airline schedules. The schedules vary on a daily basis and we consider dated time periods of one week. We consider three different schedules, from summer 2005, winter 2005 and summer 2006. For each schedule we compute solutions to the integrated aircraft routing and crew pairing problem for three different crew groups: captains, first officers and flight attendants. Each crew group has three different crew bases with different strengths at each base. Also, meal-break rules and rest requirements as well as the pay structure vary between the groups.

Each schedule contains approximately 750 flights and 22,000 connections in one week and the fleet consists of 14 aircraft. Approximately 1,200 of the connections are restricted in the sense that the sit-time exceeds the minimal sit-time by at most 30 minutes. All rules and restrictions applicable to aircraft routings and crew pairings are provided by the airline and are implemented in the algorithm. Hence, all solutions generated satisfy these rules and are ready to be used in practice.

5.2 Crew pairing

As a measure of the cost of the integrated solution we use crew pairing costs without added penalties (see objective function (6)). In all figures, LP crew pairing costs rather than IP costs are presented to obtain a more monotonic representation of the solutions. In all calculations the observed IP-gap is very small, with an average of 0.81% over all solutions generated.

Because many of the crew pairing columns generated in each solution of the iterative approach remain legal and feasible for subsequent crew pairing solutions, we can store the columns and simply recalculate their objective coefficients in each iteration. Each LP computation is “hot-started” from the optimal basis of the LP solution of the previous iteration. As a result the crew pairing solutions are obtained very efficiently during the later iterations of the algorithm since very few new columns are generated.

5.3 Aircraft routing

In the schedules considered, aircraft maintenance is usually performed at night when all aircraft are grounded at few airports. If an aircraft stays overnight at an airport with no maintenance facility, this aircraft must overnight at a maintenance base the next night. Also, there is no interaction between connections performed by the aircraft on different days. Hence, we can solve the aircraft routing problem for each day independently taking the maintenance status of the preceding night into account.

5.4 Iterative algorithm with one crew group

In Figure 2 results of the captain solutions for the schedule of summer 2005 are presented. The costs and non-robustness measures of the solution operated by the airline and the solutions generated by

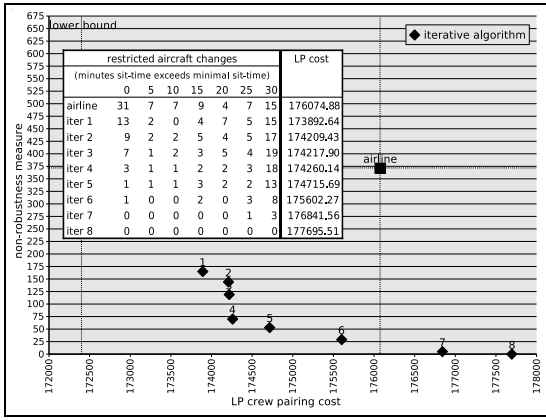


Figure 2: Comparison of airline and iterative approach, captain solutions (DPACLIM=1) for schedule summer 2005.

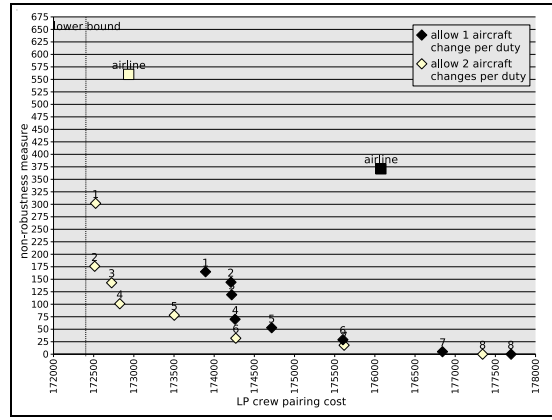


Figure 3: Variation of DPACLIM settings (black = 1; white = 2), captain solutions for schedule summer 2005.

the iterative approach are compared. The airline solution is generated by solving the crew pairing problem to cost optimality for the aircraft routing that was operated by the airline. The figure shows the non-robustness measure (see Section 4) plotted against the optimal LP crew pairing cost. In addition we list the number of restricted aircraft changes for each solution. The solutions generated by the iterative approach yield better costs and non-robustness measures than the airline solution.

The data shown in Figure 2 is given in more detail in Table 1. For each iteration we list the LP crew pairing cost and the improvement compared to the LP crew pairing cost of the airline solution in percent. We also list the IP crew pairing costs and the IP-gaps. Note that the crew pairing LP uses a weighted sum objective of crew pairing costs and penalties (see objective function (6)). The ratio of costs and penalties in the solution value of LP and corresponding IP solution can differ. When we subtract the penalties from the solution values we may obtain an IP solution that incurs less crew pairing cost than the LP solution. Finally, the non-robustness measures (NRM) and the improvements relative to the airline solution are shown. We do not show any non-robustness measure for the lower bound solution since we can not find an integrated solution satisfying the DPACLIM rule.

The cost of the crew pairing solution obtained in Step 2 of Algorithm 1 is a lower bound for the cost of the optimal solution value. In Table 1 we can observe for the captain solutions, that

Table 1: Results for technical crew, schedule summer 2005

	captain, schedule summer 2005					first officer, schedule summer 2005						
	LP cost	imp. (%)	IP cost	gap (%)	NRM imp. (%)	LP cost	imp. (%)	IP cost	gap (%)	NRM imp. (%)		
airline	176074.88	-	177647.70	0.89	371	-	172510.58	-	173176.49	0.38	344	-
lower bound	172420.84	2.08	173854.41	0.82	-	-	168419.26	2.37	170560.04	1.26	-	-
iteration 1	173892.64	1.24	175722.08	1.04	165	55.53	170079.29	1.41	171978.93	1.10	208	39.53
iteration 2	174209.43	1.06	175626.14	0.81	144	61.19	170274.45	1.30	171831.82	0.91	165	52.03
iteration 3	174217.90	1.05	175187.69	0.55	119	67.92	170544.68	1.14	171232.97	0.40	113	67.15
iteration 4	174260.14	1.03	174984.87	0.41	70	81.13	170260.76	1.30	170695.24	0.25	73	78.78
iteration 5	174715.69	0.77	175707.72	0.56	53	85.71	170388.66	1.23	170908.26	0.30	30	91.28
iteration 6	175602.27	0.27	176311.24	0.40	29	92.18	170941.18	0.91	171527.42	0.34	17	95.06
iteration 7	176841.56	-0.44	177265.20	0.24	5	98.65	172009.62	0.29	172418.49	0.24	10	97.09
iteration 8	177695.51	-0.92	178066.47	0.21	0	100.00	173395.12	-0.51	173329.55	-0.04	2	99.42
iteration 9							173909.61	-0.81	174183.32	0.16	0	100.00

Table 2: Impact of allowed aircraft changes per duty period (DPA CLIM) for captain solutions, schedule summer 2005.

	allow 1 aircraft change					allow 2 aircraft changes							
	DPAC1*	AC*	RAC*	LP cost	NRM*	DPAC1	AC	RAC	DPAC2*	AC	RAC	LP cost	NRM
airline	103	103	66	176074.88	371	76	76	55	52	104	67	172935.52	560
iteration 1	85	85	34	173892.64	165	73	73	36	25	50	27	172524.21	302
iteration 2	80	80	36	174209.43	144	59	59	20	21	42	19	172512.80	176
iteration 3	83	83	32	174217.90	119	69	69	24	16	32	12	172723.57	143
iteration 4	76	76	22	174260.14	70	58	58	18	14	28	12	172825.09	101
iteration 5	63	63	17	174715.69	53	55	55	15	8	16	8	173503.36	78
iteration 6	70	70	14	175602.27	29	66	66	8	9	18	6	174268.66	32
iteration 7	71	71	3	176841.56	5	55	55	3	14	28	4	175617.27	18
iteration 8	71	71	0	177695.51	0	60	60	0	15	30	0	177341.25	0

*DPAC1 - duty periods with one aircraft change; DPAC2 - duty periods with two aircraft changes; AC - aircraft changes; RAC - restricted aircraft changes; NRM - non-robustness measure

this lower bound solution incurs 2.08% less cost than the airline solution. The solution is infeasible since we cannot find aircraft routings to satisfy the DPA CLIM rule. Here, only one aircraft change is allowed during each duty period. The cheapest feasible solution we find (iteration 1) incurs 1.24% less cost than the airline solution and is only 0.85% more expensive than the value of the lower bound. Also, the non-robustness measure of this solution is only 165 compared to 371 for the airline solution which is an improvement of 55.53%. The most robust solution that is still cheaper than the airline solution improves the non-robustness measure by 92.18% (iteration 6). For first officers results look very similar (see Table 1). Here the non-robustness measure can be improved by 97.09% (iteration 7) without increasing the cost of the airline solution.

In Figure 3 we demonstrate the benefit of relaxing the DPA CLIM rule allowing at most two aircraft changes per duty period instead of just one. We find solutions with costs effectively on the lower bound (costs increase by 0.05% for iteration 1 and 2) that are also very robust. In Table 2 we list details for both DPA CLIM settings. We list the number of duty periods containing one (DPAC1) and two (DPAC2) aircraft changes respectively. We also list the number of aircraft changes (AC) occurring in these duty periods and how many of those are restricted aircraft changes

Table 3: Results for technical crew, schedule winter 2005

	captain, schedule winter 2005						first officer, schedule winter 2005					
	LP cost	imp.	IP cost	gap	NRM	imp.	LP cost	imp.	IP cost	gap	NRM	imp.
		(%)		(%)		(%)		(%)		(%)		(%)
airline	170767.52	-	171428.20	0.39	443	-	166361.00	-	168440.67	1.23	374	-
lower bound	168153.35	1.53	168765.37	0.36	-	-	162590.00	2.27	163235.40	0.40	-	-
iteration 1	169584.85	0.69	170965.03	0.81	267	39.73	165005.84	0.81	165997.41	0.60	199	46.79
iteration 2	169542.08	0.72	170473.83	0.55	148	66.59	164591.74	1.06	164824.08	0.14	146	60.96
iteration 3	169696.28	0.63	169922.62	0.13	121	72.69	164677.32	1.01	168909.76	2.51	97	74.06
iteration 4	169838.21	0.54	170471.04	0.37	91	79.46	165116.48	0.75	166160.70	0.63	86	77.01
iteration 5	169956.76	0.47	170514.41	0.33	76	82.84	165500.89	0.52	167020.75	0.91	65	82.62
iteration 6	171045.32	-0.16	171700.82	0.38	24	94.58	166284.73	0.05	166351.09	0.04	41	89.04
iteration 7	172390.64	-0.95	172879.06	0.28	19	95.71	166966.68	-0.36	167824.08	0.51	23	93.85
iteration 8	172787.03	-1.18	174345.83	0.89	1	99.77	168875.44	-1.51	170254.06	0.81	4	98.93
iteration 9	173217.90	-1.43	173502.53	0.16	1	99.77	169788.78	-2.06	170915.92	0.66	2	99.47
iteration 10	173201.64	-1.43	173813.00	0.35	1	99.77	170249.68	-2.34	170585.82	0.20	1	99.73

(RAC). We complement these figures with the corresponding LP crew pairing costs and non-robustness measures (NRM). We observe that the reduction in the cost objective is achieved with an increased number of duty periods with two aircraft changes. However, the non-robustness measure of these solutions does not increase significantly. In iteration 6 for example, the solution with one allowed aircraft change has a non-robustness measure of 29 and 14 restricted aircraft changes, while the solution with two allowed aircraft changes per duty period has a non-robustness measure of 32 and also 14 restricted aircraft changes. The latter solution incurs 0.76% less costs but contains 9 duty periods with two aircraft changes. This shows that we can achieve a better cost and more robust solution if we allow a small number of duty periods to contain two aircraft changes.

Table 4: Results for flight attendants, schedules summer and winter 2005

	flight attendants, schedule summer 2005						flight attendants, schedule winter 2005					
	LP cost	imp.	IP cost	gap	NRM	imp.	LP cost	imp.	IP cost	gap	NRM	imp.
		(%)		(%)		(%)		(%)		(%)		(%)
airline	68787.11	-	69833.49	1.50	497	-	68425.19	-	70046.91	2.32	511	-
lower bound	68671.77	0.17	69547.70	1.26	-	-	67909.60	0.75	70780.45	4.06	-	-
iteration 1	68749.99	0.05	70331.80	2.25	256	48.49	68142.56	0.41	70175.15	2.90	181	64.58
iteration 2	68834.40	-0.07	70285.10	2.06	129	74.04	68214.09	0.31	69025.95	1.18	149	70.84
iteration 3	68969.46	-0.27	69689.90	1.03	115	76.86	68302.00	0.18	69529.05	1.76	127	75.15
iteration 4	69161.13	-0.54	70263.65	1.57	75	84.91	68555.44	-0.19	69623.45	1.53	95	81.41
iteration 5	69473.98	-1.00	69750.55	0.40	45	90.95	68778.60	-0.52	69837.30	1.52	92	82.00
iteration 6	70230.34	-2.10	71525.90	1.81	35	92.96	70683.56	-3.30	71186.90	0.71	62	87.87
iteration 7	71926.64	-4.56	72857.50	1.28	9	98.19	71549.44	-4.57	73300.70	2.39	28	94.52
iteration 8	73313.53	-6.58	73988.75	0.91	1	99.80	74702.32	-9.17	75125.70	0.56	10	98.04
iteration 9	73289.27	-6.55	74495.55	1.62	0	100.00	74827.86	-9.36	75983.30	1.52	6	98.83
iteration 10							75073.00	-9.72	75321.95	0.33	5	99.02

Table 5: Results for technical crew and flight attendants, schedule summer 2006

	technical crew, schedule summer 2006						flight attendants, schedule summer 2006					
	LP cost	imp.	IP cost	gap	NRM	imp.	LP cost	imp.	IP cost	gap	NRM	imp.
		(%)		(%)		(%)		(%)		(%)		(%)
airline	164416.19	-	167319.42	1.74	385	-	69634.79	-	70905.51	1.79	653	-
lower bound	161398.87	1.84	161514.29	0.07	-	-	68440.59	1.71	68795.93	0.52	-	-
iteration 1	162576.29	1.12	163529.97	0.58	175	54.55	69219.33	0.60	70236.23	1.45	306	53.14
iteration 2	162769.47	1.00	165425.37	1.61	140	63.64	69376.38	0.37	70044.08	0.95	281	56.97
iteration 3	162815.10	0.97	164138.50	0.81	82	78.70	69477.70	0.23	70234.78	1.08	213	67.38
iteration 4	162896.94	0.92	163631.86	0.45	51	86.75	69458.04	0.25	69919.93	0.66	197	69.83
iteration 5	163283.65	0.69	164386.92	0.67	34	91.17	70531.49	-1.29	70393.98	-0.20	153	76.57
iteration 6	163869.72	0.33	164330.03	0.28	14	96.36	71514.31	-2.70	71535.83	0.03	121	81.47
iteration 7	163910.29	0.31	164187.37	0.17	10	97.40	73368.54	-5.36	73085.18	-0.39	84	87.14
iteration 8	165499.07	-0.66	166085.32	0.35	0	100.00	75733.24	-8.76	76139.23	0.53	53	91.88
iteration 9							76097.23	-9.28	76097.23	0.00	52	92.04

In Tables 3, 4 and 5 we present further results for technical crew and flight attendants for schedules of summer 2005 and 2006, and winter 2005 respectively. The overall trend is similar to the captain solutions for summer 2005 shown in Table 1. We find solutions that are much more robust than the airline solution and incur less or only slightly more cost. We find that for the flight

attendant solutions for summer 2006 the non-robustness measure does not decrease below 50 (see Table 5). This is due to restrictive base constraints. In experiments where the base constraints are relaxed and the number of available crew is increased by 2%, solutions with a non-robustness measure of zero can be obtained. Note also, that for the schedules of summer 2006 there exists only one set of results for technical crew since the rule sets for captains and first officers are identical for this period. Finally, the costs of the iterative solutions and of the airline solution are much closer to the lower bound for flight attendants than for technical crew. This is because of the relaxed DPACLIM rule for flight attendants where two aircraft changes in each duty period are allowed by the airline.

5.5 Iterative algorithm with multiple crew groups

We extend the iterative algorithm to solve the aircraft routing problem together with the crew pairing problems for technical crew and flight attendants. The goal is to show that the benefits of solving one crew group together with the aircraft are not lost if we add another crew group. Clearly, using a solution that is very robust for one crew group but results in many restricted aircraft changes for another crew group is not desirable. Figures 4 and 5 summarise the results. We choose the schedule of summer 2006 since this schedule only contains two different crew groups, technical crew and flight attendants. In Figure 4 the results of applying the iterative algorithm to technical crew only are shown in grey. Similarly, in Figure 5 the results of considering flight attendants only are also marked in grey. Additionally, we can solve the technical crew as before and in each iteration use the aircraft routing solution to generate solutions for the flight attendant problem. The results are shown in white in Figure 5. This corresponds to setting $w_1 = 1$ and $w_2 = 0$ in objective function (7). The solutions are more expensive and less robust than the solutions obtained from focusing on flight attendants and aircraft only. But the solutions follow the same pattern and are of good quality compared to the airline solution. Hence we do not need to sacrifice solution quality of the flight attendant problem in order to improve the technical crew solution and hence improve the overall solution to the integrated problem.

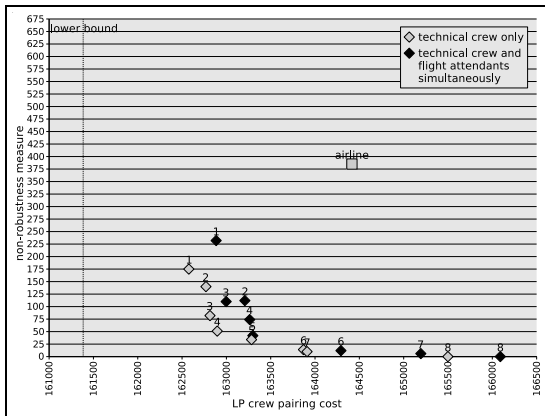


Figure 4: Results for technical crew, schedule summer 2006, with flight attendants solved simultaneously

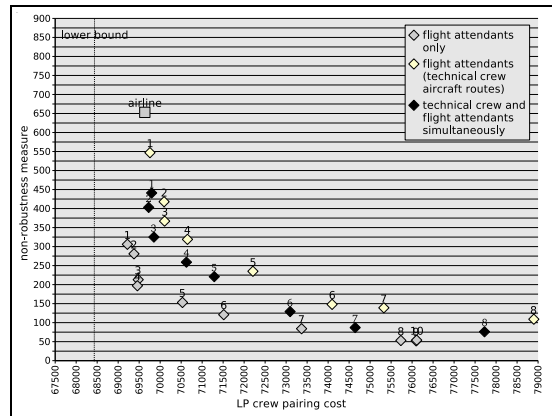


Figure 5: Results for flight attendants, schedule summer 2006, with technical crew solved simultaneously

In a second experiment we again solve two crew pairing problems in each iteration but we now use feedback information from the technical crew problem as well as the flight attendant problem to generate the aircraft routing solution. In the aircraft routing problem we weight the penalties for the restricted connections of the technical crew solution with a factor of 2 and the penalties for the restricted connections of the flight attendant solution with a factor of 1. This corresponds to setting $w_1 = 2$ and $w_2 = 1$ in objective function (7). We then generate aircraft routings subject to this objective function. The according solutions are represented in black in both figures. We observe that the technical crew solutions deteriorate slightly from the best solutions (gray). The solutions for the flight attendants are also not as good as the best flight attendant solutions (gray) but better than the solutions generated by not using any feedback from the flight attendant solutions (white). Note that in these solutions both crew groups rarely change aircraft if the connection is

restricted. This also implies that both crew groups usually stay together on restricted connections. This is of great advantage not only from a robustness point of view but also from an operational point of view. The solutions satisfy this property without the need to focus on the property in the algorithm. Since the technical crew solution usually incurs more cost than the flight attendant solution we only consider examples where the weight for the technical crew exceeds the weight for the flight attendants. The weights can be adjusted according to the airline's preferences. We can now generate solutions for the aircraft routing and two crew pairing problems in one integrated procedure. Likewise, additional crew groups or multiple aircraft types can be added.

6 Conclusion

Solving the integrated aircraft routing and crew pairing problem to optimality is computationally expensive. Also, to identify a robust solution, costs must be associated with non-robustness. In this paper we show that by heuristically coupling the two problems we can quickly generate a series of good quality solutions and do not need to attach costs to the non-robustness measure. Although optimality of the solutions can not be guaranteed, we obtain solutions that incur less costs and are significantly more robust than solutions currently used in practice. We also find that the airline can operate cheaper and more robust solutions if two aircraft changes per duty period are allowed for technical crew. Additionally, we are able to consider multiple crew groups with only a minor modification of the algorithm.

Future research includes the substitution of the crew pairing formulation with weighted sum objective by a superior bicriteria crew pairing approach (Ehrgott and Ryan [8]). We also plan to investigate solving the integrated problem to optimality and allow departure time windows for the flights to include schedule design decisions into the formulation. Departure times of the flights are allowed to vary in some window to create more flexibility for the generation of crew pairings and aircraft routings. This is expected to further improve cost and robustness of the integrated solution.

References

- [1] C. Barnhart, T. Kniker, and M. Lohatepanont. Itinerary-based airline fleet assignment. *Transportation Science*, 36:199–217, 2002.
- [2] C. Barnhart, A.M. Cohn, E.L. Johnson, D. Klabjan, G.L. Nemhauser, and P.H. Vance. Airline crew scheduling. In R.W. Hall, editor, *Handbook of Transportation Science*. Kluwer Scientific Publishers, 2003.
- [3] M.C. Bartholomew-Biggs, S.C. Parkhurst, and S.P. Wilson. Global optimization approaches to an aircraft routing problem. *European Journal of Operational Research*, 146:417–431, 2003.
- [4] L.W. Clarke, E.L. Johnson, G.L. Nemhauser, and Z. Zhu. The aircraft rotation problem. *Annals of Operations Research*, 69:33–46, 1997.
- [5] A.M. Cohn and C. Barnhart. Improving crew scheduling by incorporating key maintenance routing decisions. *Operations Research*, 51:387–396, 2003.
- [6] J.-F. Cordeau, G. Stojković, F. Soumis, and J. Desrosiers. Benders decomposition for simultaneous aircraft routing and crew scheduling. *Transportation Science*, 35:375–388, 2001.
- [7] M.S. Daskin and N.D. Panayotopoulos. A langrangian relaxation approach to assigning aircraft to routes in hub and spoke networks. *Transportation Science*, 23:91–99, 1989.
- [8] M. Ehrgott and D.M. Ryan. Constructing robust crew schedules with bicriteria optimization. *Journal of Multi-Criteria Decision Analysis*, 11:139–150, 2002.
- [9] A.T. Ernst, H. Jiang, M. Krishnamoorthy, B. Owens, and D. Sier. An annotated bibliography of personnel scheduling and rostering. *Annals of Operations Research*, 127:21–144, 2004.

- [10] T.A. Feo and J.F. Bard. Flight scheduling and maintenance base planning. *Management Science*, 35:1415–1432, 1989.
- [11] M.R. Garey and D.S. Johnson. *Computers and Intractability*. Freeman, New York, 1979.
- [12] B. Gopalakrishnan and E.L. Johnson. Airline crew scheduling: State-of-the-art. *Annals of Operations Research*, 140:305–337, 2005.
- [13] R. Gopalan and K.T. Talluri. The aircraft maintenance routing problem. *Operations Research*, 46:260–271, 1998.
- [14] M. Grönkvist. Accelerating column generation for aircraft scheduling using constraint propagation. *Computers & Operations Research*, 33:2918–2934, 2006.
- [15] C.A. Hane, C. Barnhart, E.L. Johnson, R.E. Marsten, G.L. Nemhauser, and G. Sigismondi. The fleet assignment problem: Solving a large-scale integer program. *Mathematical Programming*, 70:211–232, 1995.
- [16] D. Klabjan. Large-scale models in the airline industry. In G. Desaulniers, J. Desrosiers, and M.M. Solomon, editors, *Column Generation*, pages 163–196. Kluwer Scientific Publishers, 2005.
- [17] D. Klabjan, E.L. Johnson, G.L. Nemhauser, E. Gelman, and S. Ramaswamy. Airline crew scheduling with time windows and plane count constraints. *Transportation Science*, 36:337–348, 2002.
- [18] M. Lohatepanont and C. Barnhart. Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment. *Transportation Science*, 38:19–32, 2004.
- [19] A. Mercier and F. Soumis. An integrated aircraft routing and crew scheduling model with time windows. Technical Report G-2005-12, GERARD, 2005.
- [20] A. Mercier, J.-F. Cordeau, and F. Soumis. A computational study of benders decomposition for the integrated aircraft routing and crew scheduling problem. *Computers & Operations Research*, 32:1451–1476, 2005.
- [21] M.W. Padberg. Perfect zero-one matrices. *Mathematical Programming*, 6:180–196, 1974.
- [22] J.M. Rosenberger, A.J. Schaefer, D. Goldsman, E.L. Johnson, A.J. Kleywegt, and G.L. Nemhauser. A stochastic model of airline operations. *Transportation Science*, 36:357–377, 2002.
- [23] D.M. Ryan and B.A. Foster. An integer programming approach to scheduling. In A. Wren, editor, *Computer Scheduling of Public Transport Urban Passenger Vehicle and Crew Scheduling*, pages 269–280. North-Holland, 1981.
- [24] R. Sandhu and D. Klabjan. Integrated airline fleet and crew pairing decisions. To appear in *Operations Research*, 2006.
- [25] A. Sarac, R. Batta, and C.M. Rump. A branch-and-price approach for operational aircraft maintenance routing. *European Journal of Operational Research*, 175:1850–1869, 2006.
- [26] A.J. Schaefer, E.L. Johnson, A.J. Kleywegt, and G.L. Nemhauser. Airline crew scheduling under uncertainty. *Transportation Science*, 39:340–348, 2005.
- [27] S. Shebalov and D. Klabjan. Robust airline crew pairing: Move-up crews. *Transportation Science*, 40:300–312, 2006.
- [28] J. Yen and J. Birge. A stochastic programming approach to the airline crew scheduling problem. *Transportation Science*, 40:3–14, 2006.