J Sched (2006) 9:203–221
DOI 10.1007/s10951-006-6778-8

Operational airline reserve crew planning

Milind G. Sohoni · Ellis L. Johnson · T. Glenn Bailey

Abstract Airlines are continually faced with the challenge of efficient utilization of their cockpit crew resources. In addition to regular flying crews, some airlines have to maintain significant reserve staffing levels to meet contractual obligations and provide smooth daily operations. Reserve crews are required to cover trips remaining unassigned due to disruptions during daily operations. Airlines using a bidline system to award crew work schedules require additional reserves to cover scheduling conflicts, which result in trips dropping out of optimized bidlines. Whenever reserves are unavailable to cover these trips during daily operations, the airline has to pay a premium to cover these trips using regular pilots. The resulting operating expenses can be significant. Furthermore, inefficient utilization of reserves can cause excessive long-range crew staffing resulting in additional training and new hire expenses. In this paper, we propose a new optimization strategy to increase reserve crew utilization and build monthly reserve crew work schedules by addressing the issue of scheduling conflicts and daily operational reserve requirements.

1. Introduction

Cockpit crew manpower planning is one of the most important and challenging tasks faced by major airlines. Crew costs account for a major portion of airline operating expenses and easily

M. G. Sohoni (✉)
AC4 - Level 1, Indian School of Business, Gachibowli, Hyderabad 500019, India
e-mail: milind_sohoni@isb.edu

E. L. Johnson
School of Industrial and Systems Engineering,
Georgia Institute of Technology, Atlanta, GA 30332-0205, USA
e-mail: ejohnson@isye.gatech.edu

T. Glenn Bailey
Research, Modeling, and Design, A3 Building, 9th floor,
Delta Technology Inc., 1001 International Blvd., Atlanta, GA 30354-1801, USA
e-mail: glenn.bailey@delta.com

exceed a billion dollars annually for large airlines. Airline crews are broadly classified into two categories: *regular crews* and *reserve crews*. Regular crews are typically used to cover the flying needs and their monthly work schedules are optimized to maximize such coverage. But, in most airlines following a bidline system to award crew work schedules, a large portion of the flying drops out of these optimized work schedules due to *bidding-invoked* conflicts (explained in Section 2). Flying also remains unassigned during disruptions to normal daily operations. Reserve crews are primarily needed to cover most of this uncovered and unassigned flying.

While most airlines employ substantial resources to increase their regular crew utilization, not much effort is spent in planning for higher reserve utilization and availability. The current state-of-the-art reserve planning systems identify daily operational reserve requirements, generate legal reserve work schedules, and then use a set-covering algorithm to select a subset of reserve work schedules. A model for planning reserve work schedules at U.S Airways is described by Dillon and Kontogiorgis (1999). Gaballa (1979) describes a slightly different model to estimate reserve demand at Qantas Airways using expected overnight flight delays and call-out rates for reserve crews. Most reserve planning systems fail to capture the various sources of reserve demand, and the process is completed without much consideration to the fact that flying assigned to reserves is of varying duration. Failure to explicitly consider such constraints can cost considerable resources to reassign regular crews from other flying assignments, even though the planned reserve schedules showed ample coverage. The possibility of bidding-invoked conflicts and the lack of proper reserve planning methodologies has resulted in higher reserve staffing numbers and underutilized reserve crews. Pilot manpower has also increased though flying requirements have not proportionally increased. Some large U.S. airlines carry up to 30% reserve pilots and typical irregular operations do not justify such a high percentage.

In this paper, we propose an integrated reserve optimization strategy that increases reserve availability during operations by characterizing reserve demand in terms of *open-time* trips (explained in Section 2) and constructing reserve work schedules to cover these open-time trips. Characterizing reserve demand in terms of open-time trips allows planning for reserve demand for consecutive days rather than just daily demand. After collecting data from a major U.S. airline we also observed that scheduling conflicts with *recurrent-training* (explained in Section 2) and vacation were the two most significant sources of open-time trips. While it is difficult to control vacation assignments, scheduling recurrent-training using reserve availability constraints is possible. Hence, we also propose constructing recurrent-training schedules after the bidding-and-award process is completed. The new model controls reserve utilization by reducing open-time trips and penalizing reduction in reserve availability.

The structure of the remaining paper is as follows: Section 2 provides a brief description of the planning process and term definitions. Related literature is also reviewed in Section 2. Section 3 describes reserve schedules, reserve pay structures, and operational costs. Section 4 describes the basic algorithmic framework, the models used to estimate reserve demand, and the two phase optimization procedure to select reserve patterns. We also provide results of some related computational experiments in Section 4. In Section 5 we discuss a recurrent-training-scheduling model that explicitly includes reserve availability constraints while assigning training schedules to crew members. The computational experiments in Section 5 show the effectiveness of the training model in controlling reserve availability. Finally, we conclude the paper in Section 6.

2. The premonth planning process

Building crew work schedules includes many complex tasks. Figure 1 depicts the planning process (also called the *premonth* planning process) for constructing monthly crew work

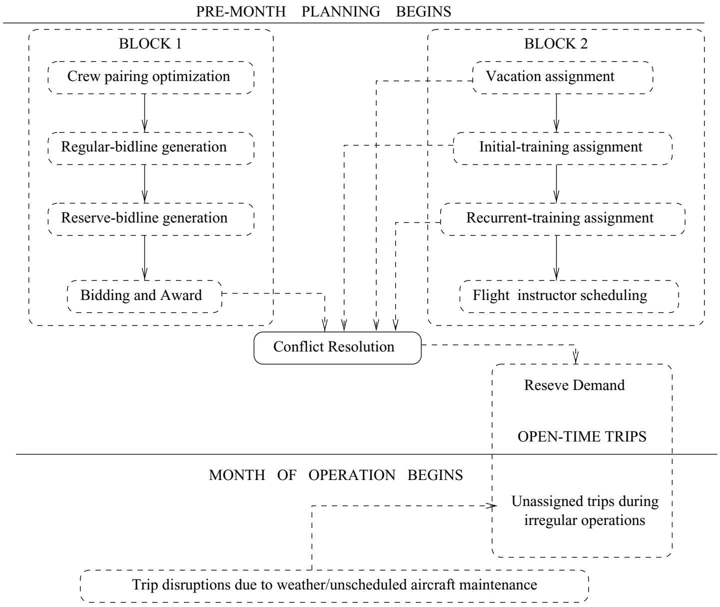

Fig. 1 The premonth planning process: Bidding-invoked conflicts result in a large number of open-time trips

schedules. Two parallel sequences of complex scheduling tasks are shown within the dotted blocks. Solid arcs indicate the sequence of planning operations and parallel blocks imply no direct dependency of processes across these blocks. Dotted arcs indicate sources contributing to scheduling conflicts, which are explained later. Schedules built during premonth planning are executed during the month of operation.

First, the crew planners build *trips* (or pairings), which are legal strings of flights, with intervening periods of rest, that begin and end at a *crew base* or *domicile* (crew pairing optimization). There is substantial literature available on optimizing trips built for regular crews (Anbil et al., 1991; Anbil, Tanga and Johnson, 1992; Barnhart et al., 2003). Barnhart et al. (2003) provide a detailed survey of the airline crew scheduling problem. Then, the crew planners construct crew work schedules, called *regular-bidlines* (or *regular-lines*) to cover these trips. Significant literature is also available on the regular-bidline generation problem (Jarrah and Diamond, 1997; Christou et al., 1999; Campbell, Durfee and Hines, 1997, Weir, 2002); though airlines tend to have a large number of business rules unique to their operation. The underlying mathematical model is a set-partitioning formulation very similar to the crew pairing optimization model. While the crew scheduling department constructs regular-bidlines with pairings, planners from the training department build and assign training schedules to eligible pilots. The training planners first build *initial-training* schedules and then *recurrent-training* schedules. *Initial-training* is a one-time event for crew members moving to a different fleet while pilots must periodically participate in *recurrent-training* to remain qualified to fly in their fleets. Initial-training affects crew manpower planning and the complexity of manpower planning and training is described in Yu, Dugan and Argüello (1997), where a heuristic approach is described to solve the problem. Qi, Bard and Yu (2004) describe an integer programming formulation for an initial-training class scheduling application at Continental Airlines. The issue of reserve crew manpower planning is not addressed specifically.

Although regular-bidlines represent work schedules, they are not personalized to meet individual pilot's assignments for training, vacation, and sick leave until the airline finalizes the bid awards. Once the crew scheduling department has constructed the regular-bidlines, crew members bid for these optimized regular-bidlines. A separate *bidding-and-award* process based on seniority ranking is used to award these regular-bidlines. During the bidding process, pilots can bid for work schedules that may conflict with their vacation or training assignments (creating *bidding-invoked conflicts*). Crew planners resolve these conflicts during a conflict-resolution phase. Since crew planners optimize regular-bidlines to maximize coverage of trips, bidding-invoked conflicts result in many trips dropping out of the regular-bidlines. These *uncovered* (dropped) trips are collectively called *open-time*; in some instances almost 40% of the trips drop out of optimized regular-bidlines into open-time. A direct conflict results in trips becoming uncovered because the training period and a scheduled trip overlap, while an indirect conflict occurs when scheduled trips and training violate regular-bidline and monthly work-legality rules. Thus, the demand for reserves is significantly increased by the need to cover open-time trips due to bidding-invoked conflicts.

Even though airlines employ crew recovery optimizers to fully utilize regular crew members during disruptions (Lettovský, Johnson and Nemhauser, 2000; Stojković and Soumis, 1998), reserve crews are still needed to provide smooth operations. These models implicitly assume that the operations control center knows about regular crew availability. This is often not the case; due to contractual guidelines. During irregular operations, if the planned crews are unable to fly their planned trips, volunteers (from the pool of regular-bidline holders) are sought to cover these trips. If no volunteers are available, reserve crew members are assigned to these trips based on availability. If reserves are not available, regular crew members may be drafted, at premium pay (typically twice the regular pay), to fly these trips. Thus, reserve unavailability is an expensive proposition, especially when voluntary flying is low.

In most bidding-and-award systems, where work schedules are not personalized, scheduling conflicts seem inevitable. This is partially due to the fact that work assignments are made independently. Personalized crew rostering systems and preferential bidding systems, adopted by some North American carriers and most European carriers, have an advantage in this respect. Gamache et al. (1998) describe a preferential bidding system application at Air Canada, and Day and Ryan (1997) describe a rostering application for cabin crews at Air New Zealand. But seniority based preferential bidding systems may suffer from the deficiency that the resulting roster may lead to perception of inequity within crew ranks. Careful adjustments have to be made to make sure that crew members understand how their preferences are scored and satisfied.

3. Reserve crew schedules

Reserve crew work schedules, unlike regular-bidlines, do not consist of pairings strung together. Instead, they consist of groups of consecutive *on-duty* and *off-duty* days, and are also referred to as *reserve patterns.* An example of a 30-day pattern is,

1 1 1 0 0 0 0 0 1 1 1 1 1 0 0 0 0 0 1 1 1 1 1 0 0 0 0 0 1 1 1 1

where, **'0'** indicates an off-duty day and **'1'** indicates an on-duty day. The total number of off-duty days and the groupings of the off-duty days determines a *pattern type.* Since pattern types depend on the contract, the potential number of legal patterns can vary from a few thousand to several million. Airlines use different types of patterns for different types of reserve crews (domestic, international) and often employ more than one type for a particular set of reserve crews. Pattern restrictions include maximum and minimum number of consecutive on-duty and off-duty days within a pattern, maximum number of *singleton off-duty* (off-duty day by itself) days, and fixed number of *golden days* (inviolable off-duty days) and their distribution within the pattern. Planning restrictions may include minimum number of off-duty reserve crews per day, and maximum allowable unavailability of reserves on weekends and holidays.

Reserve patterns are identified by the consecutive off-duty day groupings within the pattern. Thus, a 4-3-3-2 pattern type implies a pattern with exactly one consecutive off-duty grouping of length 4, exactly two off-duty groupings of length 3 and exactly one off-duty day grouping of length 2. The length of consecutive on-duty day groupings lies between 3 and 5 days. We use a similar notation to describe reserve patterns throughout this paper.

3.1. Reserve pay structure and operational costs

Regular crews and reserve crews are paid differently based on negotiated contracts. Typically, regular crew pay is a combination of the total flying hours, total duty time, and time away from base. Depending on the pilot contract, in case of a trip conflict regular crew members are paid the maximum of the trip pay or the training pay; and, if the dropped trip arises due to vacation conflict, the crew member gets paid for vacation. In either case, the airline has to pay some additional amounts to cover the dropped trip. Assuming that a trip has at least 4 hours of flying time, and that there are several hundred trips in open-time, a conservative estimate of the cost corresponding to these open-time trips is in the range of 0.5 to 1 million dollars per bid period.

By contrast reserve crews are paid for a fixed number of hours every bid period (called reserve guarantee). If a reserve flies a trip, the guaranteed hours of pay are reduced by the corresponding trip duration and the crew member is paid for the trip. Thus, reserve pay consists of adjusted reserve guarantee, and pay for flown trips, and is expressed as the sum of these two components. As described earlier, the pay incentives are different during daily operations depending on the number of volunteers and reserve availability. When we analyzed sample data from the crew scheduling department of a large airline, the premium pay range was in the range of tens of millions dollars annually. Average reserve utilization was less than 40% of the reserve guarantee in many fleets.

4. Building reserve patterns

Planning reserves to cover trips in open-time is resolved by solving the problem in two phases; the first phase generates reserve demand estimating open-time trips and the second phase described

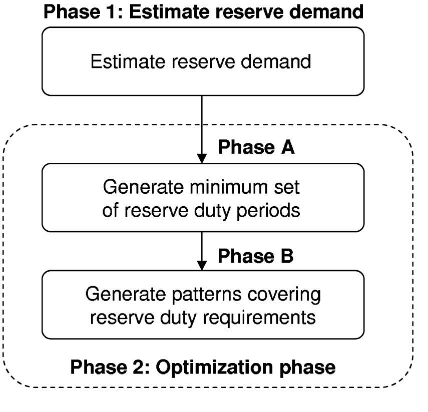

Fig. 2 The two phases of the reserve optimization algorithm: Phase 1 estimates the reserve demand and phase 2 generates patterns to cover the demand

earlier, reserve demand is due to premonth bidding-invoked conflicts and unplanned daily operations (or irregular operations) (Fig. 2). Since the open-time trip pool is used to characterize reserve demand, we use the term open-time to define trips from both these sources. We describe strategies to estimate open-time trips in Sections 4.1 and 4.2, and then we describe the optimization strategy in Section 4.3.

4.1. Estimating reserve demand due to bidding-invoked conflicts

To estimate bidding-invoked conflicts we have to understand the bidding-and-award process and pilot preferences. The *attractiveness* (C_{ij}) of a regular-bidline i to a crew member j depends on personal preferences of individual crew member j. We assume a simplistic model based on the pay incentives resulting from a conflict while ignoring any other quality of life issues. The idea is to solve a regular-bidline assignment problem with the objective of maximizing conflict. We define a binary variable, B_{ij}, such that it takes a value of 1 if regular-bidline i is assigned to crew member j and 0 otherwise. We define δ^{ij}_{cd} to be 1 if assigning regular-bidline i to crew j drops a c day long trip on day d and 0 otherwise.

We also compute the value of Γ_{cd} (from observed historical data) that limits the number of conflicting c-day long trips beginning on day d. This value of Γ_{cd} allows the predicted open-time to be realistic. Assuming that we have more crew members than regular lines, the solution to model MAXCONF

$$\text{MAXCONF} : \max \sum_i \sum_j C_{ij} B_{ij}$$

$$\sum_i B_{ij} \leq 1 \tag{4.1}$$

$$\sum_j B_{ij} = 1 \tag{4.2}$$

$$\sum_i \sum_j \delta^{ij}_{cd} B_{ij} \leq \Gamma_{cd} \forall d, c \tag{4.3}$$

$$B_{ij} \in \{0, 1\}$$

gives us the profile of the trips dropped under the strategy of creating a maximum conflict. Constraint (4.1) ensures that a crew member holds at most one regular line while constraint (4.2) ensures that a regular-bidline must be awarded to exactly one regular crew member. Constraint (4.3) restricts the model from dropping too many trips on any given day. An optimal solution to MAXCONF gives the reserve demand due to bidding-invoked conflicts.

4.2. Estimating reserve demand for daily operations

Reserve demand during daily operations is primarily due to unplanned schedule disruptions. Most real time crew recovery procedures, during disruptions, attempt to swap parts of pairings within regular crew members to provide low-cost solutions. Reserves are used if flights cannot be covered and fall out into open-time. *Move-up* crews (Shebalov and Klabjan, 2004) provide a good opportunity to swap pairings and may be described as crews that are available and legal to cover a disrupted flight sequence of another crew with the following conditions: Both crews belong to the same base, and both crews have the same number of duty hours beyond the station under consideration.

Planning for reserve crew availability, at crew bases with limited move-up crew availability, may provide additional relief without raising planned costs. Flight legs at crew bases (which

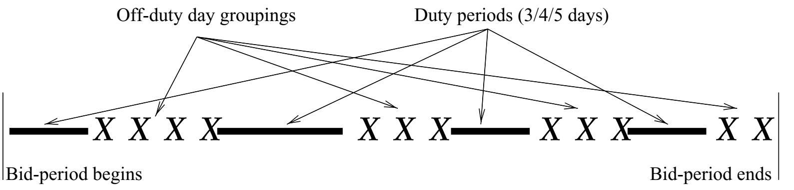

Fig. 3 A sample reserve work schedule based on a 4-3-3-2 reserve pattern

normally coincide with hubs) that do not have move-up crews are identified and the remainder portion of the trips (beyond the identified leg arriving into the base) are included with other open-time trips as potential reserve demand. In large fleets, with a large number of crews, it may suffice to include only a few trips as reserve demand. Another strategy to estimate operational reserve demand is to use a stochastic tool like *SimAir* which simulates airline operations (Rosenberger, 2001; Rosenberger et al., 2002). Given crew regular-bidlines, aircraft rotations, block times and aircraft maintenance distributions, expected weather patterns, and various recovery policies for aircraft and crews, *SimAir* simulates the operation of the flight schedule. We modified *SimAir* to simulate monthly reserve demand based on irregular daily operations (due to weather disruptions and unplanned aircraft maintenance). We integrated a crew recovery optimizer to estimate the number of disrupted trips, on each day of the bid period, that contributed to the operational reserve demand. These trips also provided an estimate for the consecutive days a reserve crew member was required.

4.3. Optimizing reserve patterns

The second phase of the algorithm generates a set of reserve patterns that cover the trips in open-time. Before we describe the optimization approach, we define the following terms: A *reserve duty period* is a string of trips which can be flown in the specified sequence without any rule violations; and, *reserve work schedule* is a bid period-long sequence of reserve duty periods mapping to a legal reserve pattern (Fig. 3).

The optimization process is split into two phases; phase A selects a set of reserve duty periods that cover all the trips in open-time, and phase B selects the required number of reserve patterns to generate reserve work schedules. The nature of trips covered in each duty period is not as critical as the number of reserve duty periods generated on each day of the month and their duration. Hence, in phase B it suffices to only deal with a smaller set of reserve duty periods that cover the open-time trips. This serves as the primary motivation to split the optimization process into two phases; the other critical reason being the ability to control the computational time to generate solutions.

The models developed for both phases are standard set-covering and set-partitioning models. The objective of phase B is to cover the maximum number of reserve duty periods with the required (or minimum) number of reserve patterns. Given a set of reserve duty periods, patterns

Table 1 Reserve pattern selection: Patterns providing higher coverage possibilities are preferred

	Patterns			Duty periods	
Day	P_1	P_2	P_3	D_1	D_2
1	0	1	1	1	0
2	1	1	1	1	1
3	1	0	1	0	1

can be selected in numerous ways to provide coverage (since each reserve pattern provides coverage for a fixed number of days). The costs associated with reserve duty periods along with the number of reserve duty periods covered by a reserve pattern provide a mechanism for distinguishing between these solutions. Consider the following example. Suppose we have a 3-day bid period with exactly two reserve duty periods D_1 and D_2 operating on days shown in Table 1. Let us also assume that we have exactly three reserve patterns P_1, P_2, and P_3 which have on-duty days (indicated by 1) as shown in Table 1. Clearly, P_1 covers D_2, P_2 covers D_1, and P_3 provides cover for both D_1 and D_2. If we have to choose exactly one reserve pattern then choosing P_3 provides a better operational solution. Since with just one reserve pattern we can cover either D_1 or D_2 during operations. Since the number of volunteers available to fly specific open-time trips are confirmed closer to the day of operation, reserve pattern P_3 provides the flexibility to cover any duty period with fewer volunteers. Similarly, if we have to choose two reserve patterns we would prefer choosing P_3 twice rather than choosing P_1 and P_2 once. Since, setting the pattern score appropriately provides good operational solutions, we employ an increasing scoring function that depends on the number of reserve duty periods covered. For our computational experiments, we use the function $I_p = \beta.\ \ln(|S_p|)$, where S_p is the set of reserve duty periods covered by pattern p and β is a positive nonzero scalar multiplier. It is assumed that $S_p \neq \emptyset$.

4.4. Phase A: Generating reserve duty periods

We assume that the length of any grouping of consecutive on-duty days within a reserve work pattern is between 3 and 5 days. This assumption is not restrictive and can be easily modified, based on business and contractual requirements, without affecting the solution procedure. Depending on the number of trips in open-time, the possible number of reserve duty periods of any duration is prohibitively large. Since generating all combinations is computationally expensive, we restrict our model to duty periods with durations between 3 and 5 days. Furthermore, the on-duty groupings within the reserve patterns determine the maximum length of a reserve duty period. The primary objective is to cover the set of trips with the minimum number of reserve

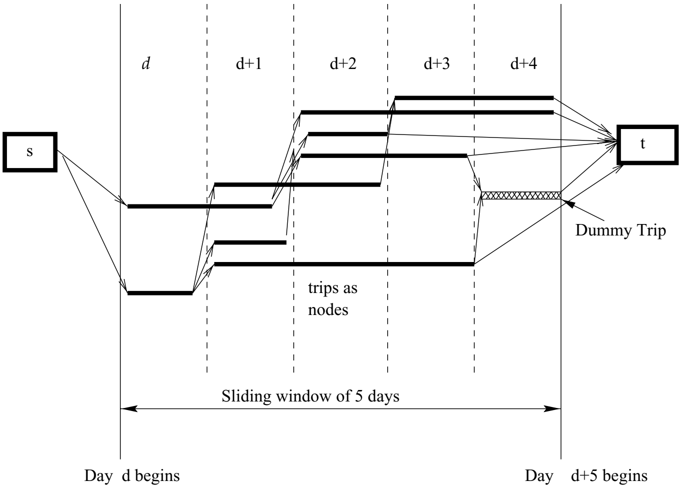

Fig. 4 Sliding window duty period graph: Reserve duty periods of the appropriate length are generated using a modified depth-first-search algorithm

duty periods of these specified durations. We try to reduce the number of reserve duty periods with *idle days* (days when the duty period does not cover any trip or a portion of a trip) by generating reserve-duty periods with the least number of idle days.

For example, consider five consecutive days starting on any day d within a month-long bid period. Construct a reserve-duty period graph, as shown in Fig. 4, where the nodes are the trips that start and end within this 5-day time window. Figure 4 shows all trips starting on day d and ending before the beginning of day $d + 5$. If no trip exists starting on a day then a *dummy trip*, one day in duration, is added at the beginning and ending on that particular day. Two special one day duration nodes, s and t, are also added to the graph. A directed arc exists from node i to node j if the starting day of j is the day after the end day of trip i; thus, we control the number of idle days in a reserve duty period. Additional directed arcs are added from node s to all nodes beginning on day d and from all nodes ending on day $d + 2$, $d + 3$, and $d + 4$ to node t. A directed path from s to t in this graph is a reserve duty period of length between 3 and 5 days. Using reserve duty period graphs for each day of the month, we generate all possible 3-, 4-, and 5-day reserve duty periods using the trips in open-time.

In general, let D define the set of duty periods generated and T denote the set of trips in open-time. Let i and j index over the sets D and T, respectively. Define C_i to be the cost of reserve duty period i and Q_j to be the cost associated with trip j. C_i depends on $\sum_{j \in i} Q_j \ \forall i$. Now, define α_{ij} such that it takes a value of 1 if duty period $i \in D$ covers trip $j \in T$ and 0 otherwise. Furthermore, define variable X_i such that it is set to 1 if duty period $i \in D$ is selected and 0 otherwise. Also define O_j such that $O_j > 0$ if trip $j \in T$ is over covered (more that one duty covers the trip). The solution to the optimization problem (MINDP), where P_j is the penalty for over covering trip j, provides a minimum cost reserve duty period cover for all the trips in open-time.

$$\begin{aligned} \text{MINDP:}\quad & \min \sum_{j \in T} P_j \cdot O_j + \sum_{i \in D} C_i \cdot X_i \\ & \sum_{i \in D} \alpha_{ij} X_i - O_j = 1 \ \forall j \in T \qquad (4.4) \\ & X_i \in \{0, 1\} \ \forall i \\ & O_j \geq 0 \text{ and integer } \ \forall j \end{aligned}$$

Constraint (4.4) enforces that each trip is covered by at least one duty period.

4.5. Phase B: Generating reserve work schedules

In most cases there are only a few thousand legal reserve patterns that meet all the contractual obligations. Hence enumerating all possible combinations of legal on-duty and off-duty days is possible. Suppose $\mathcal{J}$ represents the set of all types of reserve patterns and the set K denotes the different pattern types. Let a solution to MINDP be denoted by D^*. We need to choose, with replication, the minimum number of patterns covering the consecutive day requirements generated using D^*.

Suppose that the length of the bid period is M_1. Notice that every pattern consists of a few (typically 4 to 5) groups of consecutive on-duty days. Each grouping of consecutive on-duty days may have 3, 4, or 5 on-duty days. If a grouping of five consecutive days begins on day d, then the pattern can be used to cover any reserve duty period of lesser duration beginning on day d; but we would like to use this pattern to cover a specific duty period of a specific duration. To achieve this kind of separation we generate additional patterns, using the patterns from the base set $\mathcal{J}$, with *idle* days. Suppose the pattern is used to cover a 3-day reserve duty period beginning on day d, then days $d + 4$ and $d + 5$ are *idle* days. For every standard reserve pattern (template pattern), $p \in \mathcal{J}$, we generate all possible patterns with different on-duty days and idle day combinations while ensuring that: (i) the minimum and maximum number of consecutive working days in each group are within 3 and 5; (ii) the number of on-duty groupings are exactly the same as the base pattern p with a one-to-one correspondence between the groupings; and, (iii) the starting day of any grouping of consecutive on-duty days matches an on-duty in the template pattern p within the corresponding grouping. Generating all patterns allows the optimization model to choose patterns, built from the same template, and to target different duty periods of different lengths.

Let the patterns generated be denoted by the set $\mathcal{P}$. From reserve duty periods in D^* we compute the number of c day-long reserve duty periods beginning on day d and denote it by q_{dc}. Since reserve duties D^* contain dummy trips and over-covered trips, we split such reserve duty periods into duty periods of smaller duration that do not contain dummy trips. Thus, q_{dc} may have values for shorter duty period lengths too. This also provides the opportunity for better coverage. Let Q denote the corresponding reserve duty period requirements vector with components q_{dc}. We denote a pair of a bid period day d and c consecutive on-duty days as (d, c) and let the number of such (d, c) pairs be denoted by M_2. Clearly the dimension of Q is M_2. Note that, beginning on every day d of the bid period, a reserve pattern may have exactly c consecutive on-duty days.

Further, define $\mathcal{P}_k$ as the set of all legal patterns, in the set $\mathcal{P}$, of type $k \in K$ and let p index over the set $\mathcal{P}$. Let us also define two matrices A and B with dimensions $M_2 \times |\mathcal{P}|$ and $M_1 \times |\mathcal{P}|$, respectively. Each row in A corresponds to a (d, c) pair and each row of B corresponds to a bid-period day. The entries $a_{(d,c),n} \in A$ and $b_{mn} \in B$ are defined as follows:

$$a_{(d,c),n} = \begin{cases} 1: & \text{If pattern } n \text{ has exactly } c \text{ consecutive} \\ & \quad \text{on-duty days starting on day } d \\ 0: & \text{Otherwise} \end{cases}$$

$$b_{mn} = \begin{cases} 1: & \text{If pattern } n \text{ has an off-duty day on day} \\ & \quad m \text{ (not idle day)} \\ 0: & \text{Otherwise} \end{cases}$$

We compute the score vector, $\mathcal{I}$, with coordinates corresponding to incentive of each pattern p computed using the scoring function. Furthermore, assume that we are given the number of reserve work schedules required (R), the number of patterns of type k required (L_k), and the number of off-duty reserve crew members required on day d of the bid period (O).

If the integer variable Y_p represents the number of times a particular pattern p is selected, the optimization model, RESOPT, can be described as follows:

$$\text{RESOPT: } \min\,(P \cdot U - \mathcal{I} \cdot Y)$$

$$\sum_{p \in \mathcal{P}_k} Y_p \geq L_k \ \forall\, k \in K \tag{4.5}$$

$$A \cdot Y + U \geq Q \tag{4.6}$$

$$B \cdot Y \geq 0 \tag{4.7}$$

$$\sum_{p=1}^{|\mathcal{P}|} Y_p \leq R \tag{4.8}$$

$$U \geq 0, \text{ and } Y_p \geq 0 \text{ and integer,}$$

where P is the penalty, of dimension M_2, for not meeting the target cover value on day d for c consecutive on-duty days, and U be the corresponding under coverage vector. $Y = (Y_1, \ldots, Y_p, \ldots, Y_{|\mathcal{P}|})^{\mathrm{T}}$ represents the pattern selection vector. Constraint (4.5) guarantees the requisite number of different pattern types are selected, while constraint (4.6) ensures that the (d, c) requirements are met. Constraint (4.7) is required to guarantee that any solution has the minimum number of off-duty reserve crews on each day of the bid period. The model selects at most R patterns due to constraint (4.8).

4.6. Computational experiments

As a basis for comparison, we used a model that selects a set of reserve patterns meeting projected daily requirements. We refer to this model as MOD1. Table 2 shows a comparison between the performance of the proposed reserve model (RESOPT) and MOD1. Comparing the models allows us to demonstrate the need to characterize reserve demand in terms of both daily requirements and consecutive day requirements. The comparison is not intended to demonstrate the use of fewer reserves. The rules used to generate daily requirements for MOD1 were based on typical business rules used by the crew scheduling department. These rules compute daily reserve requirements as a percentage of the number of regular-bidlines constructed and expected number of reserves needed to cover bidding-invoked conflicts. Additionally, a fixed percentage of the total number of regular-bidlines constructed is specified as the minimum. MOD1 selects a set of reserve patterns that not only meets these daily requirements but also meets the minimum off-duty constraint.

We used 4-3-3-2 and 4-4-4 pattern types with at most five consecutive on-duty days for all the test cases. Typically the minimum off-duty constraint (4.7) guarantees 20% of the total number of reserves are off-duty. But this constraint was relaxed for these test instances. In all test problem instances, the set of template reserve patterns did not have more than 5000 patterns, i.e., $|\mathcal{J}| \leq 5000$, and $|\mathcal{P}| \leq 125000$. We used ILOG CPLEX 7.0 (see ILOG, 2000) to solve the resulting integer programming instances. All the test instances were run on an HP 9000/899 server running HP-UX 11.0. We measured the effectiveness of the proposed planning model by the number of trips left uncovered in open-time (since these trips could result in premium pay) and the minimum number of reserve patterns required. In Table 2 column (0) identifies the problem while column (1) shows the number of open-time trips of varying duration (these include dummy trips). Column (2) shows the number of 3- to 5-day reserve duty periods generated and column (3) shows the minimum number of reserve duty periods selected to cover all the trips in open-time. Column (4) differentiates the models used to obtain the solution. Column (5) shows the maximum number of reserve patterns selected in the final solution. Column (6) tabulates the overall run time for RESOPT. Column (7) shows the total number of uncovered trips in open-time. The table does not list the run time for model MOD1 since these are insignificant (less than 15 seconds). In all test instances, the Q vector computed (after removing dummy trips) from the solution to phase A (D^*) had a large proportion of short-duration reserve duties (1 or 2 days). A key observation, from Table 2, is that RESOPT leaves fewer number of uncovered trips from the open-time pool, thus controlling the expected cost exposure during daily operations. RESOPT also uses fewer reserve crews in these test instances.

4.6.1. Varying pattern types

Pattern types used within RESOPT affect the minimum number of reserves required to cover trips in open-time. Table 3 shows a comparison between five different pattern types, each with a minimum of three consecutive on-duty days and a maximum of five consecutive on-duty days

Table 2 RESOPT results: In these problem instances the model leaves fewer uncovered trips

(0) Name	(1) # Open trips	(2) # Reserve duty periods	(3) # Min. duty periods	(4) Model	(5) # Reserve	(6) CPU (seconds)	(7) # Uncovered trips
Pr_L_1	537	7219	422	MOD1	100	—	20
				RESOPT	54	385	0
Pr_L_2	448	1341	380	MOD1	60	—	19
				RESOPT	34	263	0
Pr_L_3	561	27669	367	MOD1	90	—	28
				RESOPT	52	291	0
Pr_L_4	335	933	277	MOD1	52	—	11
				RESOPT	27	262	0
Pr_L_5	350	956	282	MOD1	65	—	13
				RESOPT	34	261	0
Pr_L_6	211	397	182	MOD1	35	—	8
				RESOPT	18	282	0
Pr_L_7	183	337	154	MOD1	30	—	7
				RESOPT	17	261	0

Table 3 Reserve pattern type comparisons using the RESOPT model

		Pattern type				
		4-3-3-2	4-4-4	3-3-3-3	2-2-3-3-2	9-1-1-1
Problem	# Patterns available	(1500)	(64)	(125)	(10370)	(1288)
Pr_1	# Trips/Duties			297/837		
	# Columns	102677	13328	8613	73933	111675
	# Reserves	51	64	54	51	60
	CPU (seconds)	217	24	15	157	258
Pr_2	# Trips/Duties			258/528		
	# Columns	102677	13328	8613	73933	111675
	# Reserves	50	61	54	51	60
	CPU (seconds)	215	23	14	165	256
Pr_3	# Trips/Duties			349/4973		
	# Columns	102677	13328	8613	73933	111675
	# Reserves	51	66	58	51	60
	CPU (seconds)	222	25	16	170	262

(except near the beginning and end of the bid period). Each column represents a pattern type (denoted by the consecutive off-duty day groupings within the pattern). Three problem instances (Pr_1, Pr_2, and Pr_3), with varying open-time trips, were tested to illustrate the difference. For each problem, the number of trips in open-time, the number of reserve duty periods generated after phase A, and the total available patterns are listed. The solution shows the number of reserve patterns selected to cover the duty periods, and the computational time.

Pattern 4-3-3-2 and 2-2-3-3-2 require the least number of reserves for the obvious reason that the choice of patterns is larger in either case. Surprisingly, type 3-3-3-3 performs better than

Table 4 Varying pattern mix using the RESOPT model

Min. % 4-3-3-2	Min. % 4-4-4	# Reserves	CPU (seconds)	# Trips uncovered
20	20	52	257	0
20	30	55	257	0
20	40	58	257	0
20	50	62	304	0
20	60	65	256	0
20	70	72	257	0
0	0	51	257	0

Note. The model is sensitive to the pattern types used to solve the problem instances.

9-1-1-1, even though the number of available patterns is significantly smaller. This is due to the better distribution of on-duty days within the available patterns. The computational time is not a significant issue in either of the problems. Across the three problem instances, the variation in the number of reserve lines is not significant due to the nature of the trips in open-time.

Table 4 shows the effect of using pattern combinations within the same optimization run. The level of 4-4-4 patterns is varied from 20 to 70% while holding the level of 4-3-3-2 patterns constant at 20%. The number of reserves required to cover trips in open-time increases as the required mix of 4-4-4 patterns increases. When RESOPT is allowed to choose between 4-3-3-2 and 4-4-4 patterns (last row), the 4-3-3-2 is used exclusively in the optimal solution. All problem instances had 297 trips in open-time and 837 reserve duty periods.

5. Increasing operational reserve availability

Bidding-invoked conflicts due to recurrent-training assignments and vacation are one of the major contributors to open-time trips. We propose a scheme where recurrent-training schedules are assigned after the bidding-and-awards process is completed. Rather than assign training without knowing the trips on the regular-bidline, the idea is to build and place recurrent-training schedules on the bidline with the minimum possible conflict. While scheduling conflicts are inevitable, such a change in the process allows the recurrent-training scheduling model to control the number and duration of trips in open-time while maintaining sufficient reserve availability (even though the reserve patterns are already awarded). In Section 5.1 we briefly explain the problem, the model, and some computational results. Schoni (2002) provides a detailed description of the recurrent-training scheduling process and the algorithmic framework.

5.1. Controlling reserve availability in recurrent-training scheduling

Airline pilots must train periodically to maintain flight qualifications within a particular fleet and position (captain, first officer, or second officer). Typically more than 60% of the total monthly pilot training is due to recurrent-training and every pilot must complete training within a specified time period annually. Training includes class room instruction, called *ground school* training, and training in simulators. The total time available on simulators is divided into fixed number of time slots each day and simulator training is conducted for an entire or partial time slot depending on training requirements. Individual training requirements vary depending on current qualification level, fleet type, position, and other considerations. Pilots follow predetermined

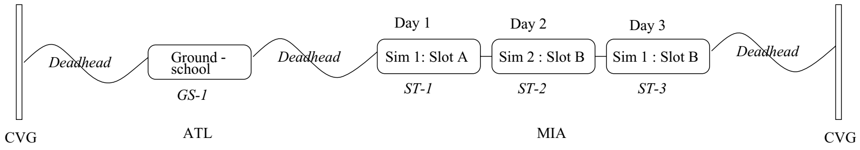

Fig. 5 Recurrent-training schedule example

training footprints which specify the training sequence. An example of a footprint designation is *4-day 767-A GS-1 ST-1 ST-2 ST-3 standard domestic* where *4-day* signifies that ground-school and simulator-training require 4 days; *A* stands for the position type (in this case a captain); *GS-1* is the code for a particular type of ground school training; and, *ST-1, ST-2*, and *ST-3* represent specific simulator training events. Figure 5 shows a training schedule of one ground school period and three simulator training periods for a 767 captain based in Cincinnati (CVG). The pilot has to deadhead to Atlanta (ATL) to attend a particular ground school training event denoted by the code *GS-1*. The pilot has to further deadhead to Miami (MIA) after ground school to attend simulator training events denoted by the codes *ST-1, ST-2*, and *ST-3*. After completing three simulator training days in MIA, the pilot deadheads back to CVG. Note that simulator training is scheduled on different physical simulators at different times during the day; this occurs due to simulator maintenance, student pairings, or simulator capabilities. Crew members typically have a fixed time window in the year to complete training. Training schedules are built for an entire bid period that is usually 30 or 31 days in duration. The bid period in which a crew member is scheduled for training, which typically corresponds to the month of hire, is called the *due* bid period and all such crews are called *due* crews. A crew member may be trained, if training capacity exists, in a bid period prior to his due bid period and such crew members are classified as *prior-to-due*. If crew members do not receive training during their prior-to-due or due bid periods, then they must receive training in the bid period that immediately follows their due bid period and are classified as *grace crews*. Grace crews have the highest priority in receiving training during any bid period since they risk losing their flight qualification the following month. Recurrent-training scheduling broadly describes the problem of finding a minimal cost pilot assignment to classroom sessions and simulator time slots, based on specific training footprints. Such assignments, including any deadhead flights if required, constitute pilot training schedules (see Fig. 5). All training, including deadheading back to the crew bases, must begin and end within the bid period.

We now present a condensed version of the model described by Schoni (2002) to illustrate the effectiveness of our proposed reserve availability constraints and objectives. Schoni et al. (2003) describe an application of a similar model at Delta Air Lines where the reserve constraints and penalties are not specifically modeled. By controlling the number of trips in open-time, the airline was able to save \$7.5 Million in operating costs. The costs associated with assigning recurrent-training to a regular crew member includes the cost of trips dropped due to conflict, training, pay, and instructor costs. The basic idea is to drop shorter trips on days when reserves are available to cover these trips. At the same time, we do not want to exhaust the reserve availability on any given day of the bid period.

Suppose that we construct a set of recurrent-training schedules, throughout the bid period, that may be assigned to individual crew members. Let i, j, and s be indices over the set of recurrent-training schedules, crew members, and simulator slots, respectively, and $c \in \{1, 2, 3, 4\}$ denote the number of consecutive days. We denote the cost of assigning recurrent-schedule i to crew

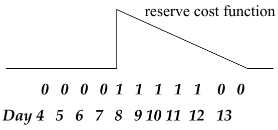

Fig. 6 A linear cost function to descriminate between on-duty days

(a) $0\ 1\ \underline{1\ 1\ 1\ 1}\ 0$ (b) $\underline{1\ 1\ 1\ 1}\ 1\ 0\ 0$

(c) $1\ \underline{0\ 0\ 0\ 1}\ 1\ 1$ (d) $1\ 1\ \underline{1\ 1\ 0\ 0}\ 0$

Fig. 7 Recurrent-training overlap with on-duty days in a pattern results in loss of reserve availability

member j by L_{ij}. In the case of a reserve crew member undergoing recurrent-training, we propose a penalty cost function corresponding to the reduction in operational availability, and overlap with off-duty days. Figure 7 shows some possibilities of recurrent-training overlap with reserve bidlines (only sections of the lines are shown). On-duty days within a reserve pattern are valued differently; Fig. 6 shows a typical cost function used to compute the cost associated with overlap. In the example shown, the cost of overlapping a training schedule with day 8 is higher than the remaining days. Assignment (c) in Fig. 7 is generally more preferable than (b). Note that the reserve availability is the least affected in (c). Similarly assignment (d) is more preferable than (a) since fewer on-duty days are violated in (d) than in (a). Furthermore, the off-duty day immediately after the end of training in (d) provides a 24-hour rest period.

We also propose including an additional penalty for dropping trips on days when reserve crews are unavailable and including reserve availability constraints for 1, 2, 3, and 4 consecutive on-duty days beginning each day of the bid period. Let this cost component be denoted by R_{cd}. Let A_{cd} denote the number of reserves available for c days starting day d, S_j denote the set of schedules for crew member j, and C_j denotes the cost of leaving crew member j unassigned. The reserve availability (A_{cd}) is computed from the reserve bidlines that have already been constructed. Further, let δ_{cd}^{ij} take a value of 1 if assigning schedule $i \in S_j$ to crew j drops a c day-long trip on day d and 0 otherwise, and α_{is}^{j} is set to 1 if schedule $i \in S_j$ uses simulator slot s and 0 otherwise.

We now define variables X_{ij}, U_j, and E_{cd}. X_{ij} is set to 1 if schedule $i \in S_j$ is assigned to crew j and 0 otherwise. U_j takes a value of 1 if crew member j remains unassigned and 0 otherwise. E_{cd} represents the amount by which reserve availability is exceeded on day d for c consecutive days.

The recurrent-training model, for assignment of a set of available recurrent-training schedules to a subset of crew members, can now be expressed as follows:

RECTRN :

$$\min \sum_j C_j U_j + \sum_d \sum_c R_{cd} E_{cd} + \sum_j \sum_{i \in S_j} L_{ij} X_{ij}$$

$$\sum_{i \in S_j} X_{ij} + U_j = 1 \ \forall_j \tag{5.1}$$

$$\sum_j \sum_{i \in S_j} \alpha_{is}^{j} X_{ij} \leq 1 \ \forall\, s \tag{5.2}$$

$$\sum_{j}\sum_{i\in S_j}\delta_{cd}^{ij}X_{ij} - E_{cd} \le A_{cd}\ \forall c, d \tag{5.3}$$

$$X_{ij}, U_j \in \{0, 1\},\ \ E_{cd} \ge 0$$

Constraint (5.1) is an assignment constraint ensuring that each crew member is assigned no more than one training schedule. Constraint (5.2) ensures a simulator slot is utilized by exactly one training schedule, while constraint (5.3) forces trips to be dropped on days reserves are available.

The solution process can be described in brief as follows. First, the process identifies all the training footprints and generates all possible legal recurrent-training schedules (using these training footprints) during the *Schedule Generation Phase*. Since all crew positions cannot be solved together, the algorithm uses a crew selection heuristic which creates subsets of crew members based on different criteria (such as position, crew base) before the *Iterative Optimization Phase*. During the *Iterative Optimization Phase*, the process iterates over these crew member subsets, and assigns schedules to maximize pairing of student pilots while minimizing the need for *filler instructors. Filler instructors* are needed when pilots cannot train in pairs (captain and first officer). At each iteration preference is given to matching a crew member's schedule to a previously assigned crew member belonging to a complementary position.

Post optimization solution improvement heuristics primarily focus on improving simulator utilization by combining schedules whenever a simulator slot is partially utilized by crew members. Partial utilization of a simulator slot is typical for some training footprints, and leaves a portion of the slot time open for other training events. The additional benefit is the reduction in the need for filler instructors.

5.2. Computational experiments

Table 5 lists real operational problems tested using the RECTRN model within the iterative optimization phase of the solution process described earlier. We solved these problem instances using the iterative algorithmic procedure described by Sohoni (2002). Column (1), in Table 5, identifies the problem and column (2) tabulates the number of crew bases. The number of

Table 5 Characteristics of the problem instances solved using the RECTRN model

(1)	(2)	(3)	(4)									(5)
			Crews									
			# Grace			# Due			# Prior			
Name	# Bases	# Simulators (slots)	A	B	C	A	B	C	A	B	C	# Training FPs
Pr_S_1	1	3 (268)	0	0	0	15	6	0	18	1	0	3
Pr_S_2	1	3 (268)	0	0	0	15	6	0	18	1	0	3
Pr_S_3	3	3 (192)	0	3	0	20	8	0	32	26	0	7
Pr_S_4	2	1 (26)	3	1	0	9	2	0	7	4	0	4
Pr_S_5	2	1 (123)	2	0	0	11	3	0	11	7	0	4
Pr_S_6	1	1 (73)	0	2	0	2	11	0	13	18	0	4
Pr_S_7	1	2 (107)	1	2	0	9	17	0	19	33	0	4
Pr_B_1	3	4 (416)	0	10	4	6	6	7	16	13	15	12
Pr_B_2	4	5 (366)	18	7	0	60	29	0	56	29	0	9
Pr_B_3	6	7 (608)	20	15	0	117	120	0	134	148	0	27

Table 6 Solution characteristics of the test problem instances solved using the RECTRN model. The loss in reserve availability is effectively controlled

(1)	(2)	(3)	(4)	(5)	(6)				(7)
					Reserve shortage				
Name	# Reserve	CPU (seconds)	% Simulator utilized	Training events	1D	2D	3D	4D	Unassigned
Pr_S_1	15	40	20	268	0	0	0	0	16
Pr_S_2	20	39	20	268	0	0	3	0	16
Pr_S_3	22	62	47	192	0	0	5	0	28
Pr_S_4	30	10	100	26	0	0	0	0	10
Pr_S_5	35	27	35	123	0	0	0	2	11
Pr_S_6	40	12	42	73	0	0	0	0	13
Pr_S_7	40	22	55	107	0	0	0	0	21
Pr_B_1	60	325	17	416	0	0	0	0	17
Pr_B_2	80	738	60	366	0	0	2	4	58
Pr_B_3	80	794	67	608	0	0	1	3	249

simulator and ground-school locations are one for all problem instances. Column (3) shows the number of simulators and simulators slots available while column (4) shows the number of crews classified by their training status and position (captain, first officer, or second officer). Column (5) shows the total number of unique training footprints for each problem.

We used ILOG CPLEX 7.0 (ILOG, 2000) to solve these specific MIP instances on an HP-9000/200 server. Table 6 shows the solution characteristics using this algorithm. The maximum run time was restricted to 1000 seconds for each sub grouping of pilots, thus ensuring run times are within acceptable real time planning limits. The restricted problem has typically less than 80,000 columns per iteration. In most cases the gap between the LP objective value and IP objective value for this restricted problem was less than 0.5%. Column (2) shows the number of reserve crews available to cover dropped trips and column (3) shows the CPU time in seconds for an entire run that includes all the optimization iterations. Column (4) shows the percentage of simulators utilized after recurrent-training scheduling is completed. Simulator utilization is not 100% in most cases since the remaining capacity is leased out or used for initial training. Furthermore, the restrictions on the number of positions per base also affects simulator utilization. Column (5) shows the total number of simulator-training-events schedules, which in turn determines the number instructors required, while column (6) shows the number of one, two, three and four day trips remaining uncovered after recurrent-training is completed (based on a certain reserve availability profile). Column (7) shows the number of crew members remaining unassigned at the end of the scheduling process. Crew members remain unassigned primarily due to capacity restrictions. In almost all cases, the reserve constraints effectively control the number of trips left uncovered. The number of reserve patterns for problems Pr_S_1 through Pr_S_7 ranged between 10 and 40. Problems Pr_B_1 through Pr_B_3 had 60 to 80 reserve patterns available. In most instances the model was also able to control the length of the trips dropped, thus minimizing conflict costs.

6. Conclusions

Reserve crew costs are a major portion of operating budgets for most major airlines and efficient utilization of reserves is necessary to control overall crew costs while ensuring smooth daily operations. As seen by empirical test results, RESOPT effectively increases reserve availability. Combined with a good model to estimate trips in open-time, RESOPT can also help control reserve manpower, thus lowering long range manpower costs significantly. Even though it is extremely difficult to estimate open-time reserve demand due to irregular operations, adding a reserve set for known open-time trips (due to conflicts) to the aggregate needs for irregular operations (computed using traditional methods) will result in a significant improvement in reserve availability. The MAXCONF model needs to be improved further to reflect an appropriate objective function.

As shown in Section 4.6.1, different reserve patterns result in different reserve coverage. While the number of available legal reserve patterns may vary significantly across the types, the effect of using these reserve patterns may not be very significant. Additional steps to reduce trips in open-time (Section 5.1) help in further increasing reserve crew availability and utilization.

A related area of interest is estimating the number for *short call* and *long call* reserves (a definition based on the time allowed before reporting to duty on the day of operation) on each day of operation. Simulating daily airline operations could forecast the number of expected trip disruptions every hour and the number of short (and long) call reserves. Though such a computation could be useful for recovery, the benefits and risks need to be carefully assessed.

References

Anbil, R., E. Gelman, B. Patty, and R. Tanga, "Recent advances in crew-pairing optimization at American airlines," *Interfaces*, **21**, 62–74 (1991).

Anbil, R., R. Tanga and E. L. Johnson, "A global approach to crew-pairing optimization," *IBM System Journal*, **31**, 71–78 (1992).

Barnhart, C., A. Cohn, E. Johnson, D. Klabjan, G. Nemhauser, and P. Vance, "Airline crew scheduling" in R. W. Hall (ed.), *Handbook of Transportation Science*, Kluwer Scientific Publishers, Boston, Massachusetts, pp. 493–521 (2003).

Campbell, K. W., R. B. Durfee, and G. S. Hines, "FedEx generates bidlines using simulated annealing," *Interfaces*, **27**, 1–16 (1997).

Christou, I., A. Zakarian, J. Liu, and H. Carter, "A two phase genetic algorithm for large scale bidline generation problems," *Interfaces*, **29**, 51–65 (1999).

Day P. R., and D. M. Ryan, "Flight attendant rostering for short-haul airline operations," *Operations Research*, **45**, 649–661 (1997).

Dillon, J. E. and S. Kontogiorgis, "US airways optimizes the scheduling of reserve flight crews," *Interfaces*, **29**, 123–131 (1999).

Gaballa, A., "Planning callout reserve for aircraft delays," *Interfaces*, **9**, 78–86 (1979).

Gamache, M., F. Soumis, D. Villeneuve, J. Desrosiers, and E. Gelinas, "The preferential bidding system at air Canada," *Transportation Science*, **32**, 246–255 (1998).

ILOG, "*ILOG CPLEX 7.0 User's Manual*," ILOG, France, 2000.

Jarrah, A. and J. Diamond, "The problem of generating crew bidlines," *Interfaces*, **27**, 49–64 (1997).

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

Qi, X., J. F. Bard, and G. Yu, "Class scheduling for pilot training," *Operations Research*, **52**, 148–162 (2004).

Rosenberger, J. M., 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).

Rosenberger, J. M., *Topics in Airline Operations*, Ph.D. thesis, Georgia Institute of Technology, Atlanta, GA, 2001.

Shebalov S., and D. Klabjan, "Robust airline crew pairing : Move-up crews" Working paper, The University of Illinois at Urbana-Champaign, 2004. Available from http://netfiles. uiuc.edu/klabjan/www.

Sohoni, M. G., A Robust Optimization Approach to Reserve Crew Manpower Planning in Airlines, Ph.D. thesis, Georgia Institute of Technology, Atlanta, GA, (2002).

Sohoni, M., G. Bailey, K. Martin, H. Carter, and E. Johnson, "Delta optimizes continuing-qualification-training schedules for pilots," *Interfaces*, **33**, 57–70 (2003).

Stojković, M. and F. Soumis, "The operational airline crew scheduling problem," *Transportation Science*, **32**, 232–245 (1998).

Weir, J. D., A Three Phase Approach to Solving the Bidline Generation Problem with an Emphasis on Mitigating Pilot Fatigue through Circadian Rule Enforcement, Ph.D. thesis, Georgia Institute of Technology, Atlanta, GA, 2002.

Yu, G., S. Dugan, and M. Argüello, "Moving toward an integrated decision support system for manpower planning at Continental airlines: Optimization of pilot training and assignments," in G. Yu (ed.), *Industrial Applications of Combinatorial Optimization*, Kluwer Academic Publishing, Boston, Massachusetts 1997.