

Large-scale Models in the Airline Industry

Diego Klabjan

Department of Mechanical and Industrial Engineering
University of Illinois at Urbana-Champaign, Urbana, IL
klabjan@uiuc.edu

December 8, 2003

Abstract

Operations research models are widely used in the airline industry. By using sophisticated optimization models and algorithms many airlines were able to improve profitability. In this paper we review these models and the underlying solution methodologies. We focus on models involving strategic business processes as well as operational processes. The former models include schedule design and fleetings, aircraft routing, and crew scheduling, while the latter models cope with irregular operations.

1 Introduction

In the United States the Airline Deregulation Act of 1978 gave the airlines much more commercial freedom to compete. Since then, to leverage demand with capacity or sit inventory, the airlines have pioneered revenue management. Among other breakthroughs, to offer a variety of itineraries, major airlines have developed the so-called *hub-and-spoke* networks. On the other hand, to improve profitability they use sophisticated tools for reducing cost. In recent years, the raise of low-fare, no-frill airlines such as Southwest in the U.S. and Ryanair in Europe put additional pressure on the remaining carriers. To keep low fares, the airlines must maintain low cost per airline-sit-mile. This is commonly achievable through contract renegotiations and by using enhanced modeling and optimization techniques.

Since the 1950s the airlines are using operations research models in solving their complex planning and operational problems. These models have become increasingly complex. On the one hand, the airlines have become larger (through mergers or expending service) resulting into large-scale models. On the other hand, the continuing pressure to increase profitability resulted into more “accurate” models and better solution methodologies. For example, an access crew cost of several percent was acceptable a decade ago but it is not today. In many cases, by using state-of-the-art crew scheduling decision support systems the access cost has been pushed below one percent. Large-scale models have become computationally tractable due to algorithm, hardware, and software advances.

On the algorithmic front the most notable advance has been the introduction of *column generation*. In column generation, a model is given implicitly and is dynamically updated in order to improve the incumbent solution. Such an approach enables handling of the entire large-scale problem and at the same time it reduces the computational burden.

In this paper we review large-scale linear models that are frequently encountered and used in the airline industry. We also outline in Section 2 the underlying methodologies for solving these models. We start in Section 3, by explaining business processes in airline planning and operations. In Section 4 we present models that concern the passenger service. Models for schedule planning and fleetings are given in Section 4.1, then we review aircraft scheduling in Section 4.2, and at the end we discuss crew scheduling in Section 4.3. For every problem we discuss planning and operational models. Recent trends are presented in Section 5.

2 Solution Methodologies for Large-scale Models

Here we briefly overview three most common techniques for solving large-scale linear mixed integer models: branch-and-price, Lagrangian decomposition, and Benders decomposition. We start with branch-and-price.

Large-scale linear programs are often solved by *delayed column generation*. In this algorithm, at every iteration, only a subset of columns is considered. The problem with only a subset of columns is called the restricted master problem. In every iteration of the algorithm, first the restricted master problem is solved and let π be the optimal dual vector, which for ease of discuss we assume it exists. Next the so called *subproblem* is solved. In subproblem solving we identify a set S of columns with the lowest reduced cost with respect to π . If we cannot find a column with negative reduced cost, then we stop since π is an optimal dual solution to the original problem and together with the optimal primal solution to the restricted master problem we have an optimal primal/dual pair. Otherwise, we append columns in S to the restricted master problem and the entire procedure is iterated. When the restricted master problem includes too many columns after several iterations, columns with large reduced cost are removed from the restricted master problem.

Frequently the most computationally intensive step in delayed column generation is subproblem solving since it needs to scan many columns and typically it is a complex task to generate a single one. When columns correspond to constrained paths in a network, an efficient algorithm known as *constrained shortest path* is often employed, [Desrosiers et al. \(1995\)](#), [Desaulniers et al. \(1998\)](#). In this case, the task is to find the cheapest cost $s-t$ path (reduced cost in delayed column generation framework) among all paths with certain properties. We explain the algorithm by an example. Assume we want to find a shortest path in a network subject to the duration and the number of arcs in paths being below a given number. By duration we mean that every arc has an associated transit time and the duration of a path is the sum of transit times along the path. We introduce label vectors, which in this case have 3 coordinates. The first one corresponds to the cost, the second one to the duration, and the last one to the number of arcs. With every node we associate a set of label vectors. For example, a label vector $(-45, 134, 4)$ at node i corresponds to an $s-i$ path with cost -45 , duration 134, and 4 arcs. The constrained shortest path algorithm uses the same framework as standard shortest path algorithms. Suppose the algorithm selects a node i for scanning. The constrained shortest path algorithm next scans all neighbors j of i and all label vectors $k = (k_1, k_2, k_3)$ at i . Each label vector k is updated by traversing the arc (i, j) and the updated label vector is appended to node j . In our example, the label update means that the new label vector has $k_3 + 1$ as the third component, k_2 plus the transit time of arc (i, j) as the second component, and k_1 plus the cost of arc (i, j) as the first component. The key observation is that under some realistic assumptions label vectors that are dominated can be discarded. If we have two label vectors \bar{k}, \tilde{k} at node j and $\bar{k} \leq \tilde{k}$ component-wise, then the $s-j$ path corresponding to \bar{k} is not going to be part of the shortest path. The efficiency of the algorithm depends heavily on the frequency at which dominance occurs. Note that if there is no dominance, the algorithm simply enumerates all paths. It turns out that dominance occurs often in practice and therefore the algorithm is computationally efficient. Two alternative algorithms for subproblem solving in presence of constrained shortest paths are sketched in [Section 4.3](#).

Branch-and-price is a branch-and-bound algorithm, where LP relaxations at every node are solved by delayed column generation. Since subproblems are often combinatorial in nature, the standard variable dichotomy is not appropriate. When columns correspond to constrained paths in a network, the following branching strategy is frequently used. Let r, s be two adjacent nodes, which are selected based on the incumbent LP solution. Then in one branch only paths where node s immediately follows node r are considered. This is easily reflected in the network by removing all arcs from r , except the (r, s) arc. On the other branch we forbid all paths where s does follow r , which is captured by removing the (r, s) arc. This branching rule produces more balanced tree. Since LP relaxations tend to be computationally intensive, only few branch-and-bound nodes are evaluated. For this reason, a common strategy is to use depth-first search and abort after the first integer solution is obtained. An excellent survey on branch-and-price is given by [Barnhart et al. \(1998a\)](#).

Another common technique is *Lagrangian decomposition*, [Fisher \(1981\)](#), [Fisher \(1985\)](#), [Martin \(1999\)](#). Suppose we can partition constraints into “easy” and “difficult” constraints. The concept behind is that if

the difficult constraints are removed, the resulting problem is easily solvable. In Lagrangian decomposition, every difficult constraint gets a linear penalty and it is moved to the objective function. The resulting problem is called the Lagrangian relaxation and it is a function of the penalties. Let us assume that we have a maximization problem. For any given values of penalties, the Lagrangian relaxation is computationally easy. It is easy to see that it always provide an upper bound on the optimal solution. The goal now is to find the best upper bound, i.e. to minimized the Lagrangian relaxation over all possible penalties. This is the Lagrangian dual problem, which is a nonlinear optimization problem. In practice it is solved by variants of the subgradient algorithm. One drawback of this approach is that there is no guarantee to find feasible solutions. They have to be constructed heuristically during the execution of the subgradient algorithm. The algorithm is very appealing since it is easy to implement and it handles complex side (difficult) constraints.

The *Benders decomposition*, [Benders \(1962\)](#), [Minoux \(1986\)](#), is well suited for mixed integer programs with linking integer variables. It requires that for any fixed value of integer variables, the resulting problem is an LP, where the constraint matrix is often block diagonal. The algorithm at every iteration solves a mixed integer program (restricted master problem) with a single continuous variable that provides a bound on the optimal solution. Next the linear program resulting from the original problem by fixing integer variables to the values from the restricted master problem is solved. The optimal dual vector to this LP provides a Benders cut, which is added to the restricted master problem and the procedure is repeated. The same framework can be used for convex problems.

Another technique, called constraint programming, is paving its way to the large-scale linear programming. In this manuscript we do not discuss constraint programming and relevant literature.

3 Airline Planning and Day of Operations

In this section we review typical business processes used by combination airlines, [Figure 1](#). While every airline has its own processes and its own organization names, most of the airlines follow the depicted processes and terminology. The time frames can vary significantly.

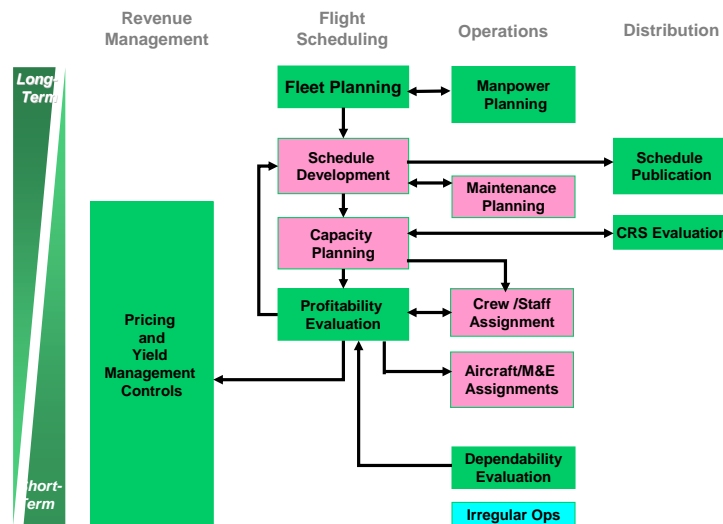


Figure 1: Business processes

Long term fleet and manpower planning consists of making strategic decisions with respect to the number of aircraft and the fleet decomposition, and cockpit crew manpower planning. In fleet planning considerations such as the airline’s mission (e.g., Southwest has a single fleet that allows relatively simple and efficient operations), aircraft utilization, route structure, cargo/passenger mix, etc., are taken into account.

The schedule development phase typically starts 12 months before the day of operations and it lasts up to 9 months. In the first phase the airline establishes the *service plan*, which is the set of services to operate in a given market. The service plan is either daily for domestic operations and weekly for international, long-haul service. The marketing group considers several factors such as traffic forecasts, status of competing carriers, internal resources, and marketing initiatives. Marketing initiatives are approved by upper management and involve decisions such as entering a new market. The designed service plan typically does not divert substantially from the current schedule. Following the service plan, the scheduling group generates a detailed *flight schedule*, i.e. a flight departure and arrival time. The flight schedule has to obey a set of operating constraints, e.g. maintenance planning, and given generic resources such as the number of aircraft. The schedule is then published. Next is capacity planning or *fleeting*. In *fleeting* an equipment type is assigned to each flight subject to available resources such as the number of aircraft. The goal of the *fleet assignment model* is to maximize profit. The schedule together with the sit capacity is then input to the computer reservation system (CRS). The produced *fleeting* solution is then evaluated with a profitability evaluation model and potential improvements are fed back to the schedule development group for possible minor adjustments. The schedule development phase and *fleeting* are discussed in Section 4.1.

Once the equipment types are assigned, *aircraft routing* and *crew scheduling* follow. In aircraft routing, called also maintenance routing, which is discussed in Section 4.2, a specific tail number is assigned to each flight subject to maintenance constraints. The objectives are usually incentives such as throughs and robustness. The goal of crew scheduling, see Section 4.3, is to assign crew members to individual flights in order to minimize the crew cost and maximize various objectives related to contractual obligations, quality of life, and crew satisfaction. Crew schedules have to satisfy complex regulatory and contractual rules. Potentially crew planners detect unfavorable connections and give feedback to schedule and fleet planners. The crew scheduling process typically starts three months before the day of operation and it is constantly updated until a few weeks before the day of operations.

Only minor changes to *fleeting*, aircraft routes and crew schedules are made during the last few weeks before the day of operations. To better match demand with capacity, some airlines perform dynamic fleet and aircraft swaps, known also as demand driven dispatch or D³ for short, [Berge and Hopperstad \(1993\)](#), [Talluri \(1996\)](#), [Clark \(2000\)](#). If preferential bidding is used, approximately one month before the day of operations, crews bid for their monthly crew assignments and only minor changes such as two way trip swaps are performed in the last few weeks.

Throughout the strategic planning processes pricing and yield or revenue management are actively involved. In revenue management, the airline controls the sit inventory by adjusting fare prices, setting overbooking limits, and making decisions at any given time about selling particular fare classes on a given passenger itinerary. Since models and solution methodologies in revenue management and pricing are substantially different than the remaining models resulting from the aforementioned processes, they are not discussed here (see e.g. [van Ryzin and Talluri \(2002\)](#) for a survey on revenue management). We also do not discuss the cargo side of planning and operations.

The actual day of operations, called also *execution scheduling*, consists of making final minor adjustments to the flight schedule (e.g., adjust arrival times based on the daily wind forecast), executing the pre-planned schedule (e.g., file the flight plan) and rescheduling for irregular operations or *disruption management*. The latter is carried out by operations controllers, which are typically located in the airline operations control center (see e.g. [Clarke et al. \(2002\)](#)). Most frequent sources of irregular operations are weather, unscheduled maintenance, congestion, crew unavailability, security problems, etc. Disruption management is composed of three processes. When an irregular operation occurs, first the aircraft are rerouted, which is called *aircraft recovery*. In this stage in addition to rerouting the aircraft, decisions on delaying and canceling flights are made. Next is the *crew recovery process*, where crews are assigned new crew itineraries. The controllers can use original, standby, and reserve crews. At the end is the *passenger reaccommodation process*, where passengers are rerouted to alternative itineraries. Clearly the new schedule must conform to all regulatory and contractual rules. While the airlines often impose more stringent rules in planning, in operations they typically use precise rules. Contractual rules for operations are usually different from those in planning.

4 Models for Passenger Service

In this section we focus on the passenger side of planning and operations.

4.1 Schedule Planning and Fleeting

For most of the airlines schedule planning is a manual process mostly driven by marketing requirements. On the other hand, decision support tools for fleeting are common. There are only few manuscripts on schedule planning but there is vast literature on fleeting. Since research papers that address schedule planning, cover fleeting as well, we start with fleeting.

The basic *fleet assignment model* (FAM), called also the leg-based fleet assignment model, is to find an optimal assignment of equipment types to flights. The input consists of a list of flights, which are given by the destination/origin station and departure/arrival time, a set of equipment types and the corresponding number of aircraft for each equipment type. Since each equipment type has its own sit capacity, on a given flight the equipment type decision can produce low load factor (lost revenue of using too large sit inventory) or a potential spill of passengers to competitors if the realized demand is higher than the sit capacity of the assigned equipment type. The typical objective function consists of the variable and fixed cost of operating a flight by a given equipment type and an estimate of potential revenue.

Next we formally describe the FAM, see e.g. [Abara \(1989\)](#), [Hane et al. \(1995\)](#). First we define the *flight time-space network*. The network has a node (u, i) for each time when an arrival or departure of leg i occurs at station u . If an event corresponds to a departure, then let t_i be the departure time of flight i . If it corresponds to an arrival, then t_i is the arrival time of flight i plus the minimum plane turn time (the so-called ready time). We assume that the activity times t_i are ordered in time, i.e. $t_1 \leq t_2 \leq t_3 \cdots \leq t_l$, where l is the number of activities at the station. There is a *flight arc* $\{(u, i), (v, j)\}$ for each leg that departs at station u at time t_i and arrives at station v at time t_j . In addition there are *ground arcs* $\{(u, i), (u, i+1)\}$ for each u and i , where we assume that we have a wraparound arc between the last and first node of the day. Each station has exactly one wraparound arc.

The model has two types of variables, the fleet assignment variables x and the ground arc variables y . For each leg i and for each fleet k there is a binary variable x_{ik} , which is 1 if and only if leg i is assigned to fleet k . For each ground arc g and for each fleet k we define a nonnegative variable y_{gk} that counts the number of planes in fleet k on the ground in the time interval corresponding to g . Let et be a fixed time typically corresponding to a time with low activity, e.g. 3 am. The FAM model reads

$$\begin{aligned} \min \sum_{\substack{i \in A \\ k \in K}} c_{ik} x_{ik} \\ \sum_{k \in K} x_{ik} = 1 \quad i \in A \end{aligned} \tag{1}$$

$$\sum_{i \in O(v)} x_{ik} - \sum_{i \in I(v)} x_{ik} + y_{o(v)k} - y_{i(v)k} = 0 \quad v \in V, k \in K \tag{2}$$

$$\sum_{g \in W} y_{gk} + \sum_{i \in M} x_{ik} \leq b_k \quad k \in K \tag{3}$$

$$y \geq 0, \quad x \text{ binary,}$$

where

$I(v)$: set of flight arcs to node v	A : set of all flight arcs
M : set of flights in the air at et	K : set of all fleets
$O(v)$: set of flight arcs from node v	V : set of nodes
b_k : number of aircraft in fleet k	$i(v)$: ground arc to node v
W : set of ground arcs containing et	$o(v)$: ground arc from node v
c_{ik} : cost of assigning fleet k to leg i .	

Constraints (1) require that each leg is assigned to exactly one fleet, (2) express the flow conservation of aircraft, and (3) assure that we do not use more aircraft than there are in a fleet.

This basic FAM model is relatively easily solvable even for large flight networks by commercial integer programming solvers, Hane *et al.* (1995). The model can be enhanced by incorporating some aircraft maintenance and crew requirements, Clarke *et al.* (1996), Rushmeier and Kontogiorgis (1997), explicitly modeling aircraft routes, Barnhart *et al.* (1998b), and incorporating departure time decisions, Rexing *et al.* (2000), Desaulniers *et al.* (1997), Bélanger *et al.* (2003). The biggest drawback of this model is the revenue component of the objective function. In a multi-leg passenger itinerary a capacity decision on a flight effects the number of passengers spilled from the itinerary and therefore the revenue contribution of other flights. Kniker (1998) explores several alternatives to compute the cost component c but none of them captures *network effects* accurately. Therefore the model has to be augmented to capture multi-leg passenger itineraries.

For ease of discussion we assume that every passenger itinerary has a single fare and we assume that passengers are not recaptured, i.e. if the booking demand exceeds the sit inventory on a given flight, the non-booked passengers are not captured on airline's alternative itineraries. Under these assumptions we next present the *passenger mix model*, which decides how many booked passengers to have in any itinerary given a fixed leg sit inventory. Let P be the set of all itineraries. The fare of itinerary $p \in P$ is denoted by f_p and let C_i be the available sit inventory of leg i . Let $w_p, p \in P$ be the decision variable that counts the number of booked passengers on itinerary p . The model for the optimal number of booked passenger reads

$$\begin{aligned} \max \quad & \sum_{p \in P} f_p w_p \\ \sum_{i \in p} w_p & \leq C_i \quad \text{for every leg } i \end{aligned} \tag{4}$$

$$\begin{aligned} w_p & \leq D_p \quad p \in P \\ w & \text{ integer.} \end{aligned} \tag{5}$$

Here $i \in p$ represents that leg i is part of itinerary p . For every $p \in P$, the unconstrained demand is denoted by D_p and it can be obtained either by a direct O-D (origin-destination) forecasting method or by segregating leg based demand forecasts. Constraints (4) impose the sit capacity limits and (5) meet the forecasted demand. Enhancements and generalizations of this model are given in Kniker (1998).

A fleet assignment model that captures O-D itineraries, called the *origin-destination fleet assignment model*, is obtained by combining the leg based FAM and the passenger mix model. The only required modification is to replace the right-hand side of (4) by $\sum_{k \in K} \tilde{C}_k x_{ik}$, where \tilde{C}_k is the sit capacity of equipment type k . While the leg based FAM is relatively easily solvable, this is not the case for the O-D fleet assignment model. The number of variables and therefore constraints in this model can be as high as 200,000. (Note that in the presence of multiple fare classes per itinerary, (5) are no longer the simple upper bound constraints.)

Barnhart *et al.* (2002b), Kniker (1998) solve the model by branch-and-price. The pricing step is not computationally intensive and it is done by a simple scan routine, i.e. there is no need for constrained shortest path. They report computational times of several hours just to find the first integer solution. Indeed, even solving the LP relaxation of the model takes 2 hours and half for a realistic model consisting of 70,000 itineraries and approximately 2,000 legs. The authors enhance the solution methodology by employing sophisticated preprocessing techniques and valid inequalities. In order to improve tractability, Barnhart *et al.*

(2002a) develop an alternative model. Instead of having decision variables that assign single legs to a fleet, the new model requires decision variables that assign a subset of legs to a fleet. Thus the assignment flight leg variables are grouped together. Clearly considering all possible subsets of legs is not tractable, however, the authors show that by carefully selecting subsets, the resulting model is tractable. Another alternative formulation to O-D fleetings is presented in [Jacobs *et al.* \(1999\)](#), where the underlying model is solved by Benders decomposition.

Next we discuss models that incorporate schedule design decisions. Until recently, algorithms for schedule design and fleetings were mostly iterative in nature, [Berge \(1994\)](#), [Marsten *et al.* \(1996\)](#), [Etschmaier and Mathaisel \(1985\)](#). Given a schedule, first demands are estimated by a schedule evaluation model. Next the FAM is solved by using the computed demands and the resulting solution is evaluated. In order to modify the schedule, flights for addition or deletion are identified. The profit resulting from addition or deletion of these flights is then estimated and based on the resulting profit a subset of these flights is selected for addition and deletion. These flights are then added to the schedule and the procedure is repeated.

Recently models that consider fleetings or aircraft routes, and schedule design decisions simultaneously emerged. Most of them are still iterative as they dynamically generate passenger itineraries and evaluate schedules, however, given a subset of itineraries, the decision of which flights to use and the fleetings decision are made simultaneously. [Lettovský *et al.* \(1999\)](#) give a model that can construct a schedule from scratch. As part of the input are service frequencies, origin-point of presence, and demand information. The model then generates a schedule that maximizes revenue subject to basic operational constraints. Their objective function is nonlinear since they use a logit-based market-share model. They solve the model by Benders decomposition (no details are given in their publication). [Lohatepanont and Barnhart \(2002\)](#) present a linear model that given a set of mandatory and optional flights, selects a subset of optional flights that maximizes total revenue. They use the O-D fleetings model and augment it with optional flights. The nonlinear relation in flight demands is taken into account by solving several models iteratively and adjusting demands based on the incumbent solution. In each iteration the model is solved by branch-and-price. [Yan and Tseng \(2002\)](#) and [Yan and Wang \(2001\)](#) present a similar model but they solve it by Lagrangian decomposition. They relax all constraints except flow conservation constraints of passenger and aircraft. [Erdmann *et al.* \(1999\)](#) present an approach for scheduling flights of charter carriers. They go a step further since they explicitly model aircraft routes. They solve the model by branch-and-price.

[Antes \(1997\)](#) presents common business processes used by the airlines in schedule planning. [Berdy \(2002\)](#) gives an excellent review on nuts-and-bolts of route generation. These two manuscripts do not present mathematical models.

4.2 Aircraft Routing

4.2.1 The Planning Stage

In tactical planning after each flight has an assigned equipment type, the aircraft routing problem follows. In this stage each individual aircraft or *tail number* is assigned to each flight in a given time period. Note that the fleetings solution decomposes the flight schedule and therefore there is an aircraft routing problem for every fleet (e.g. Boeing 737-300 and 737-400 fleets yield two separate routing problems).

In addition to the assignment requirement that each flight must be assigned a unique tail number, the routes should not use more than the available number of aircraft and they must meet maintenance requirements. In the U.S., the FAA requires four types of checks. The A-checks or line maintenance are routine checks (visual inspection of major systems), which have to be performed approximately every 65 block hours and a certain number of take-offs. Durations of A-checks are typically from 3 to 10 hours and they are usually performed during the night. B-checks are typically done once in several months and they require detailed visual inspection. For C- and D-checks an aircraft is taken out of service for a month and they are done once every one to four years. Since these two checks are spaced at large intervals, they do not pose scheduling difficulties. For this reason aircraft routing solutions consider only A- and B-checks. Maintenance checks can only be performed at specific maintenance stations, which are typically separate for each fleet. In order to decrease unscheduled maintenance events, many airlines impose more stringent

maintenance requirements, e.g. A-checks every 40 block hours and even frequent more stringent checks. In addition to these regulatory maintenance rules, some airlines impose equal utilization of aircraft, called also the big cycle constraint.

It is extremely difficult to assign a single cost attribute to an aircraft routing solution. Some airlines consider the routing problem as a pure feasibility problem. Often a value of a routing solution is a weighted sum of several attributes such as the contribution from throughs (the benefit of offering certain non-stop connections) and robustness measures to possibly decrease occurrences of unexpected events.

In the planning stage, typically several weeks or months in advance, first generic aircraft routes are constructed during a rolling time horizon. This is the *aircraft rotation problem*. These generic routes satisfy short maintenance requirements such as A-checks but do not consider, for example, B-checks and aircraft positions at the beginning of the time horizon. Only a few weeks or even days before the day of operations, the actual tail numbers are assigned to each flights, i.e. the *aircraft assignment problem* is solved. The assignment follows generic routes as much as possible but it takes into account longer maintenance requirements and the actual aircraft position at the beginning of the horizon.

Both problems are modeled either as a multicommodity flow problem or a partitioning/packing problem. Next we present a partitioning formulation from [Barnhart *et al.* \(1998b\)](#) for the rotation problem, which assumes a rolling time horizon and only checks that have to be done periodically and the period is shorter than the time horizon.

Suppose we are given a flight schedule (of a single equipment type) in a time horizon. A *string* is an ordered sequence of flights that originates and terminates at a maintenance station. The arrival station of a flight in a string is equal to the departure station of the next flight in the string and the connection times are longer than or equal to the minimum plane turn time. In addition, a string is maintenance feasible, e.g. the sum of the block times of the flights in the string is less than the one imposed by A-checks and the number of flights in a string is less than the maximum number of takeoffs between two A-checks. An *augmented string* is a string with the maintenance time interval attached to the end of it. The maintenance is assumed to start as soon as possible and it lasts for the duration of the required check. For example, if an aircraft arrives at 4pm local time and the maintenance cannot start before 8pm, it is assumed that the maintenance indeed starts at 8pm and lasts for the required period. Let S be the set of all augmented strings. A decision variable x_s is 1 if augmented string $s \in S$ is in an aircraft route and 0 otherwise. To combine augmented strings together, we need ground arc variables at maintenance stations MS , which is similar to the FAM. As in [Section 4.1](#), we define ground arcs $y_{i(v)}, y_{o(v)}$ for $v \in V$. V corresponds to activities at stations in MS and ready time is defined based on the termination time of augmented strings. The model reads

$$\begin{aligned} \min \sum_{s \in S} c_s x_s \\ \sum_{i \in s} x_s = 1 \quad \text{for every flight } i \end{aligned} \quad (6)$$

$$\sum_{\substack{j \in O(v) \\ j \in s}} x_s - \sum_{\substack{j \in I(v) \\ j \in s}} x_s + y_{o(v)} - y_{i(v)} = 0 \quad v \in V \quad (7)$$

$$\sum_{g \in W} y_g + \sum_{s \in S} r_s x_s \leq b \quad (8)$$

$$y \geq 0, \quad x \text{ binary.}$$

Here c_s is the cost of augmented string s , b is the number of aircraft in the fleet, and r_s counts how many times augmented string s crosses time et , where et is defined as in the FAM. Constraints (6) require that each flight be assigned to a string, flow balance at maintenance stations is guaranteed by (7), and (8) is the plane count constraint. Note that due to the flow balance constraints, strings can always be concatenated into an aircraft rotation. The big cycle constraint can be modeled in the similar way as the subtour elimination constraints in the traveling salesman problem (see [Barnhart *et al.* \(1998b\)](#) for details). Additional constraints such as capacities at the maintenance stations can easily be embedded. [Barnhart *et al.* \(1998b\)](#) solve this model by branch-and-price. The subproblem is solved by the constrained shortest path algorithm. For every

maintenance requirement there is a label, i.e. we must maintain a label for block hours and number of takeoffs, and we must use labels for any nonlinear cost component. If the big cycle constraint is imposed, then row generation is required as well since this implies an exponential number of additional constraints.

Cordeau *et al.* (2001) and Mercier *et al.* (2003) model the aircraft assignment problem as a multicommodity network flow with nonlinear resource constraints. The resource constraints model maintenance requirements. The model is solved by a combination of Benders decomposition and branch-and-price. The pricing problem is solved by the constrained shortest path algorithm. Sriram and Haghani (2003) use the multi-commodity formulation as well. They model maintenance requirements as linear constraints and therefore their formulation is very complex. The solution methodology is a heuristic based on local search.

Clarke *et al.* (1997) consider the aircraft rotation problem. They modeled it as an Eulerian tour with side constraints. The side constraints capture maintenance requirements. Since the Eulerian tour problem is equivalent to the traveling salesman problem on the line graph, they actually solve the traveling salesman problem. This transformation enables them to capture the big cycle constraint as the subtour elimination constraints. The model is solved by Lagrangian decomposition, where the maintenance and subtour constraints are relaxed. The underlying master problem then becomes a simple assignment problem.

Feo and Bard (1989) and Daskin and Panayoyopoulos (1989) model the assignment problem as the set partitioning problem. In such a formulation each aircraft route corresponds to a column in the formulation. The former work solves the underlying model heuristically. They first generate a set of routes for each aircraft independently. Next they solve the resulting partitioning problem by a greedy heuristic to obtain the solution. Daskin and Panayoyopoulos (1989) rewrite the formulation as a set packing model. One family of constraints require that each flight is in a route and the other one that each route is selected at most ones. They solve the model by Lagrangian decomposition, where the latter constraints are relaxed.

Paoletti *et al.* (2000) give details on aircraft rotation and assignment at Alitalia. The rotation problem is solved as an assignment problem, where maintenance requirements are not considered. They maximize the throughs value and the aircraft turn times. The assignment problem is solved a day before the day of operations and is considerably more complex. It tries to follow the solution from the assignment problem as much as possible. Their model is string based but it has several additional operational constraints. They employ a constraint programming approach.

A completely different framework is given by Gopalan and Talluri (1998) and Talluri (1996). They approach the rotation problem from a combinatorial point of view. They model the problem as the Eulerian tour problem. The former work considers 3 day maintenance checks and they show that if only these checks are required, the problem is polynomially solvable. The 4 day checks are addressed in the latter manuscript. In this case the problem becomes NP-hard and they propose several heuristics. In both cases the maintenance requirement means that an Eulerian tour must visit certain nodes (maintenance stations) every 3 or 4 arcs since their arcs correspond to lines of flying (day's activity of an aircraft).

4.2.2 Day of Operations

In this section we cover the execution part of aircraft routing. In a day of operations, due to unexpected events such as inclement weather or unscheduled maintenance, new aircraft routes have to be found.

As is the case in the planning stage, two types of models are found: the multicommodity ones and set partitioning models. The solution methodologies are either local search techniques or integer programming heuristics.

Early work on aircraft recovery is presented in Jarrah *et al.* (1993). They model the recovery problem on a time-space network. They consider cancellations and delays separately, i.e. for each one of them they have a different model. The underlying network is a pure minimum cost network optimization model and thus it does not include any side constraints. Yan and Young (1996) and Yan and Lin (1997) consider delays, cancellations, and aircraft ferrying in a single multicommodity flow model with side constraints. They solve the model by Lagrangian decomposition. A quadratic programming formulation is presented by Cao and Kanafani (1997a,b). The underlying model is a multicommodity flow model with side constraints, however side constraints are moved to the objective function with quadratic penalty terms. Thengvall *et al.* (2000, 2003) present a multicommodity network flow model with side constraints. In the former, they solve the

model with a commercial integer programming software while in the second they apply the bundle algorithm after relaxing the flight covering constraints. The former work introduces a new objective of deviation from the original schedule and they consider only minor disruptions. In most of the instances the LP relaxation gives an integer solution and if this were not the case, they use rounding to obtain a feasible solution. The delays are modeled by introducing several copies of a single flight, each one with a different departure time. [Bard *et al.* \(2001\)](#) present a similar model but they focus more on airport closures or reduced slot capacity (e.g., when the ground delay program is in effect).

Multicommodity flow models are not appropriate for capturing maintenance requirements and therefore they are suitable for small to medium disruptions. In large disruptions it takes longer to return back to normal operations and therefore maintenance constraints become an issue. Partitioning formulations, where variables correspond to complete routes, are used instead. [Løve *et al.* \(2002\)](#) present a local search heuristic approach for solving the problem. They minimize total delay, number of cancellations, and the number of aircraft swaps. [Argüello *et al.* \(1997a,b\)](#) present the underlying set partitioning formulation, which is then solved by the greedy randomized adaptive search procedure. [Rosenberger *et al.* \(2001\)](#) give a similar set partitioning formulation. Their decision variables correspond to flight cancellations and they have binary variables that assign an aircraft route to a specific tail number. The basic constraints are to assign each aircraft to a route (ferrying, diversions, and over-flying are allowed) and that each flight must be either covered by a route or cancelled. They model slot availability as well. The solution methodology consists of first selecting a subset of routes and then finding a solution over these routes by means of a commercial integer programming solver.

4.3 Crew Scheduling

4.3.1 The Planning Stage

In tactical planning, after the aircraft routes are obtained, crew scheduling is next. Crew scheduling itself is decomposed into two processes.

In the crew pairing phase crew pairings or itineraries are obtained. A *pairing* is a sequence of flights, where the destination station of a flight in the sequence corresponds to the origin station of the next flight. In addition, the origin station of the first flight and the destination station of the last flight must correspond to the same crew base. In the crew pairing stage, a pairing is not assigned to a particular crew member. The crew pairing problem is to find a least cost subset of pairings that partition all flights.

After pairings are obtained for a given time period (typically a month), individual crew members are assigned to these pairings. *Rostering* is a common process outside of North America. Given crew preferences for individual pairings and patterns, an assignment of pairings to crew members is sought in rostering. The objective consists of meeting as many preferences as possible and to minimize potential costs. *Preferential bidding* is commonly used by North American carriers. This process consists of first generating bidlines (generic monthly assignments) and then crew members based on seniority bid for bidlines.

Crew Pairing

A *duty* is a subsequence of a pairing that comprises a working day of a crew. Connection times within a duty, called sit connections, are short (35 minutes to a few hours) whereas connection times between duties, called layovers or rests, are much longer (10 hours and more). A pairing must satisfy many regulatory rules. To name just a few of them, there is a minimum sit and rest time, the elapsed time and the flying time of a duty is upper bounded, and there is the complicated 8-in-24 rule imposed by the FAA. In addition to these rules, union rules complicate pairing structure even further (maximum number of days in a pairing, more complex duty elapsed times). On top of all this, the cost of a pairing is complex. Often the cost of a pairing is the maximum of three quantities: a fraction of the pairing elapsed time, sum of duty costs in the pairing, and the number of duties times the minimum guaranteed pay. Linear terms capture hotel and meal expenses. The cost of a duty is the maximum of three terms as well: a fraction of the flying time, a fraction of the elapsed time, and the minimum guaranteed pay. Some airlines offer a fixed salary to crews and therefore their objective is to minimize the number of crews.

Most often the problem is modeled as the set partition problem with side constraints. Let P be the set of all pairings and for a $p \in P$ let c_p be the cost of pairing p . The model reads

$$\begin{aligned} \min \sum_{p \in P} c_p x_p \\ \sum_{i \in p} x_p = 1 \quad \text{for every leg } i \\ x \text{ binary,} \end{aligned}$$

where x_p is 1 if pairing p is selected and 0 otherwise. In practice side constraints are added, which most often model equal use of resources. For example, if at crew base cb there are only a given number of crews, then

$$l_{cb} \leq \sum_{p \in S_{cb}} x_p \leq u_{cb}$$

is added, where S_{cb} is the set of all pairings starting at crew base cb and l_{cb}, u_{cb} are the lower, upper bound on the number of available crews at cb , respectively. Other typical side constraints are to balance pairings across crew bases with respect to cost, the number of days of pairings, or the number of duties.

This problem is computationally challenging for the following two reasons. Each pairing has complex feasibility rules and cost structure. In addition, the number of pairings even for a medium size problem is enormous. Fleets with 200 flights can have billions of pairings. For this reason, whenever there is a repetition of flights in the time horizon, the crew pairing optimization is performed in three steps. In the first step the so called *daily problem* is solved. This is the crew problem solved over a single day time horizon and it is assumed that every flight is operated every day. Once a daily solution is repeated over the real time horizon, some pairings become infeasible (called broken pairings). The operational legs of these pairings are then considered in the *weekly exceptions problem*, Barnhart *et al.* (1996). The final solution then consists of daily pairings without the broken pairings and the pairings from the weekly exceptions problem. The weekly exceptions problem is a special case of the so called *weekly problem*, where pairings from the end of the horizon wrap around to the beginning of the horizon. The main distinction between a daily problem and a weekly problem is that in the former problem a pairing cannot cover the same leg more than once while this is allowed in the latter problem. When transitioning from one (monthly) work schedule to another, the *dated problem* needs to be solved to account for pairings that span both months. In the dated problem, flights on specific dates are given and they have to be partitioned by pairings.

A standard approach is to view pairings as constrained paths in either the *flight network*, Minoux (1984), Desrosiers *et al.* (1991), or the *duty period network*, Lavoie *et al.* (1988), Anbil *et al.* (1994), Vance *et al.* (1997). The flight network has a node associated with each departure and arrival. There is a flight arc connecting each departure node of a flight with the arrival node of the same flight. In addition, there are connection arcs between any two arrival and departure nodes with the arrival station of the first flight being equal to the departure station of the second flight and the connection time is within legal limits, i.e. the time is either between the minimum and maximum sit connection time or between the minimum and maximum rest time. In addition, the network has two artificial nodes s and t . Node s is connected to every departure node of a flight that can start a pairing. Similarly, every arrival node of a flight that can end a pairing is connected to node t . Every pairing is an $s - t$ path in the flight network. Due to various pairing feasibility rules that cannot be embedded in the flight network, every $s - t$ path is not necessarily a pairing. The duty period network is constructed in a similar way except that flight arcs are replaced by duty periods and connection arcs correspond to legal rest connections. It is assumed that duties are enumerated beforehand. The duty period network captures more feasibility rules since all duty legality rules are embedded in the network, however, it requires much more storage.

The literature on crew pairing optimization is abundant with Barnhart *et al.* (2003) providing more details and surveying the literature. Here we focus only on branch-and-price related aspects and we survey only branch-and-price related literature. In branch-and-price type algorithms subproblem solving is done on either of the two networks. There are three approaches to find a low reduced cost pairing (subproblem solving). The first one, pioneered by Desrosiers *et al.* (1991), is by constrained shortest path, Lavoie *et al.*

(1988), Anbil *et al.* (1994), Vance *et al.* (1997). In this approach a label is maintained for every feasibility rule that is not embedded in the network, e.g. 8-in-24 rule, elapsed time rules, etc. In addition, if the cost of a pairing is nonlinear, then each component in the maximum needs to have a separate label. If the duty period network is used, fewer labels are needed. For U.S. domestic carriers, the number of labels on the flight network can be as high as 20. A second approach is used in the commercial crew pairing solver from Carmen Systems, Galia and Hjorring (2003). Their approach is based on finding the k th shortest path. They find a shortest path on the current network. If the path is not feasible, they modify the network so that the obtained path is no longer a path in the network. Once a feasible path is found, it corresponds to a k th shortest path in the original network for a k . The third approach is to perform a depth-first search enumeration of pairings on a network, Marsten (1994), Anbil *et al.* (1998), Klabjan *et al.* (2001b), Makri and Klabjan (2003). Since there are too many pairings, the search has to be truncated by, for example, not considering all the duties and all connection arcs. Another enhancement is by pruning the search earlier due to some lower bounds on the reduced cost, Anbil *et al.* (1998), Makri and Klabjan (2003).

Crew pairing branch-and-price algorithms employed tailored branching rules, which are based on the branching rule designed for set partitioning, Ryan and Foster (1981). The most widely used rule is to branch on *follow-ons*. In this branching rule, two flights r, s are selected and branching follows the scheme presented in Section 2. Follow-on branching is used in Anbil *et al.* (1998), Desrosiers *et al.* (1991), Vance *et al.* (1997), and Anbil *et al.* (1994). An alternative branching rule, called *timeline branching* is proposed in Klabjan *et al.* (2001b). In timeline branching two flights r, s are selected and a connection time t . In one branch the rule requires that only pairings with the connection time between r and s less than t are considered and the other branch considers only pairings with the connection time larger than or equal to t .

Reserve crew planning and training scheduling is discussed in Sohoni and Johnson (2002a,b), and Sohoni *et al.* (2003). Klabjan *et al.* (2001a) present a model and solution methodology to solve the weekly crew pairing problem that is not based on the traditional daily/weekly exceptions paradigm. All of the related material presented so far relates to scheduling of cockpit crews. The flight attendant problem, Day and Ryan (1997), Kwok and Wu (1996), is similar except that several flight attendants are required to cover a flight. These problems tend to be larger since the flight attendants are cross qualified but, on the other hand, the feasibility rules are computationally easier.

Rostering and Preferential Bidding

Once a set of pairings is obtained that covers all flights in a month, these pairings and additional tasks such as reserve crew duties and flight training, are next assigned to individual crews. The problem decomposes further, not only based on the equipment type, but also based by the crew member rank (such as Captain, First Officer).

Feasibility rules in rostering are even more complicated than the pairing feasibility rules. The rules are imposed either by a regulatory agency such as the FAA, the airline itself, and there are contractual rules. Some of the basic requirements are: limits on the rest time between two tasks, limits on a working period (working week) between 4 to 8 days, limits on the number of monthly and yearly block hours. Then there are restrictions with respect to task coverage, e.g. one captain and one first officer for a given task, two captains and one first officer for simulator training, etc. Rules involving several rosters are common as well, e.g. some crew members prefer to fly together (married couples) and language restrictions.

In rostering several objectives are possible. From the airline perspective, minimizing open time or unassigned activities is important. Open time consists of tasks that are not assigned to regular crews but they are covered either by reserve crews or overtime is used. Clearly these two options are costly to the airline. If the number of block hours of a crew member is larger than a certain limit, the airline has to additionally pay the crew member for the overtime. Therefore the airline's interest is in minimizing the overtime pay. The third objective of the airline is to optimize assignments to training on simulators. These type of training is mandatory and very expensive. The airline also tries to produce rosters that are equitable across crew members, e.g. the crew members should have equal flying time and the number of off days. On the other hand, crew members have their own goals and preferences. Each member has its own preferences such as starting duties early in the morning, favoring certain pairings, etc. A quality roster must meet as many

preferences as possible. Additional details on rostering rules and examples are given in [Kohl and Karisch \(2003\)](#).

The rostering problem can be modeled in the following way, [Gamache and Soumis \(1998\)](#), [Gamache et al. \(1999\)](#), [Kohl and Karisch \(2003\)](#). Let the decision variable x_s^k be one if roster s is selected for crew member k . The model reads

$$\min c_s^k x_s^k \tag{9a}$$

$$\sum_{\substack{k \in K \\ i \in s}} x_s^k \geq n_i \quad \text{for every task } i \tag{9b}$$

$$\sum_s x_s^k = 1 \quad \text{for every crew member } k \tag{9c}$$

$$x \text{ binary,} \tag{9d}$$

where c_s^k is the cost of assigning roster s to crew member k , and n_i is the number of crew members that are required for task i . (9b) guarantee task coverage and (9c) assign a roster to every crew member. Rules that involve several rosters have to be explicitly modeled by adding side constraints.

Similarly to the crew pairing approach, there exists an underlying network such that a roster is a path in this network but not necessarily the other way around, [Gamache et al. \(1998\)](#). To exploit individual preferences, it is actually convenient to construct a network for every individual crew member. This leads to branch-and-price approaches. [Gamache et al. \(1998\)](#) and [Gamache and Soumis \(1998\)](#) are the first ones to describe a branch-and-price algorithm. An important observation from their work is that it is beneficial to construct crew rosters for individual crew members that are disjoint with respect to tasks. Subproblem solving is performed by constrained shortest path. [Kohl and Karisch \(2003\)](#), and [Kharraziha et al. \(2003\)](#) use k th shortest path in subproblem solving and they present a general modeling language to capture feasibility rules and objectives. The same modeling language is used also in their crew pairing optimizer, [Galia and Hjorring \(2003\)](#). To warm start the algorithm, they construct rosters heuristically. [Day and Ryan \(1997\)](#) describe cabin crew rostering at Air New Zealand in their short-haul operations. The problem is solved by first assigning off days and it is followed by assigning pairings and other tasks. This two phase approach simplifies the problem but it can lead to suboptimal solutions. In each phase they employ branch-and-price.

Many pure heuristic approaches to rostering have been developed by various airlines. They can be found in various proceedings of [The Airline Group of the International Federation of Operational Research Societies \(AGIFORS\)](#) meetings. A detailed description of a simulated annealing heuristic approach is given in [Lučić and Teodorović \(1999\)](#).

There are two approaches to preferential bidding. The first one is essentially identical to rostering except that personal preferences of crew members are not considered. A different approach is given in [Gamache et al. \(1998\)](#). Their methodology consists of producing individual rosters sequentially one by one in a given order (e.g. seniority). Suppose rosters for the first $k - 1$ crew members have already been obtained. The roster for the k th crew members is obtained by solving (9) with the following changes. The objective function considers only rosters of crew member k . (9b) are included only for those tasks that are not covered by the first $k - 1$ crew members. There is (9c) constraint for every non assigned crew member. Clearly only rosters feasible to unassigned crew members are considered. The model produces an optimal roster for crew member k and at the same time it guarantees a feasible solution for the remaining unassigned members. If there are m crew members, then m models are solved. To improve the execution time, each model is relaxed to allowing fractional solutions to rosters of crew members $k + 1, k + 2, \dots, m$. They further improve the algorithm by adding cuts. [Achour et al. \(2003\)](#) enhance this work by combining branching decisions and cuts.

[Jarrah and Diamond \(1997\)](#) present a heuristic approach to preferential bidding. A crew planner sets parameters, e.g. the length of a working period, number of off days. If the parameters are restrictive enough, there are not many rosters to consider and the resulting set partitioning model is solved by explicitly enumerating all rosters. If there are too many rosters to consider, they employ a local search heuristic. A simulated annealing heuristic to preferential bidding is given in [Campbell et al. \(1997\)](#).

4.3.2 Day of Operations

In disruption management, the crew recovery problem follows aircraft recovery. The input to the problem are the new aircraft routes together with the new departure times and flight cancellations. In crew recovery new crew assignments have to be obtained. Depending on the airline, non disrupted crew members can be involved in the reassignment or not. But clearly the number of such crew member should be minimized. Another objective is to return back to the original crew schedule as soon as possible. Then there is the objective of minimizing the cost, which can consist of the direct salary based cost, uncovered flight cost, crew deadheading, etc. In crew recovery standby and reserve crews can be used but the latter are costly.

Teodorović and Stojković (1990) develop a sequential approach based on a dynamic programming algorithm, using the first-in-first-out principle to minimize the crews' ground time. Wei and Yu (1997) present a heuristic-based framework for crew recovery. Song *et al.* (1998) present a multicommodity integer network flow model and a heuristic search algorithm to solve it. Stojković *et al.* (1998) present a column generation approach similar to the one used for crew pairing problems. Lettovský *et al.* (2000) and Lettovský (1997) base their column generation approach on the rostering model. They give details on how to quickly generate promising pairings. Stojković and Soumis (2001) incorporate flight scheduling decisions into the crew recovery nonlinear multicommodity flow model. Together with new crew assignments, their model produces new departure times. The model is solved by branch-and-price, where the subproblem is solved by constrained shortest path. Stojković and Soumis (2003) expand this model by allowing crews to split, i.e., if a first officer and a captain in planning are assigned to cover a given flight, the recovered schedule might keep the first officer at the same flight but is assigned a different captain.

5 Recent Advances

In recent years models and optimization based methodologies that integrate the three planning areas started to emerge. Integration of aircraft routing and crew pairing is discussed in Cohn and Barnhart (2002), Klabjan *et al.* (2002), Cordeau *et al.* (2001), Mercier *et al.* (2003). Solving the combined fleet and aircraft routing model, Barnhart *et al.* (1998b), has already been discussed in this manuscript. Barnhart *et al.* (1998c) take the first step towards a model for integrating fleet and crew pairing. All of this integration efforts are in an early stage and most of the methodologies are not yet suited for large-scale problems. Another obstacle in adopting these models by the airlines is that they require changes in business processes. Legacy carriers are notorious for their unwillingness to change their internal processes. On the other hand, smaller, mostly low-cost carriers are more flexible and open to business process reengineering, Garvin (2000). This fact goes hand in hand with the current inability of solving large-scale integrated models. Clearly the airlines have to follow and embrace the advances in modeling and algorithms, and the researchers have to improve decision support systems to be more tractable.

The other emerging trend is in robustness. It is well documented that customer complaints, delays, and flight cancellations were on a rise every year from 1996 until 2000. They reached the top in summer 2000, where it even caught attention by the Congress. In early 2001 tactical models that embed robustness emerged. These models do not necessarily produce a cost/profit optimal solution but a suboptimal solution that fares better in operations under uncertainty. On the crew pairing side, approaches by Schaefer *et al.* (2000), Ehrgott and Ryan (2001), Yen and Birge (2000), and Chebalov and Klabjan (2002) provide robust solutions. Robust fleet solutions are discussed in Rosenberger *et al.* (2003), Kang and Clarke (2003), and Listes and Dekker (2003). Ageeva (2000) presents an approach to robust airline routing. A robust approach to passengers rerouting in disruption management is given by Karow (2003). While many sources of frequent disruptions (congestion being the dominant one) have abated since the events of September 11, new ones are popping up (increased security measures). Nevertheless, delays have been drastically reduced due to a substantially lower demand and therefore the airlines have lost poise for robust solutions. However, the airline industry is recovering and not far in the future the demand will be at the pre September 11 level. So even though robust solutions have lost appeal in the industry, the researchers are seeing this direction as the next big step in improving profitability and customer satisfaction.

References

- Abara, J. (1989). Applying integer linear programming to the fleet assignment problem. *Interfaces*, **19**, 20–38.
- Achour, H., Gamache, M., and Soumis, F. (2003). Branch and cut at the subproblem level in a column generation approach: Application to the airline industry. Technical Report G-2003-34, GERAD.
- Ageeva, Y. (2000). *Approaches to incorporating robustness into airline scheduling*. Master’s thesis, Massachusetts Institute of Technology.
- Anbil, R., Barnhart, C., Johnson, E., and Hatay, L. (1994). A column generation technique for the long-haul crew assignment problem. In T. Ciriani and R. Leachman, editors, *Optimization in Industry II*, pages 7–24. John Wiley & Sons.
- Anbil, R., Forrest, J., and Pulleyblank, W. (1998). Column generation and the airline crew pairing problem. In *Extra Volume Proceedings ICM*. Available from <http://www.math.uiuc.edu/documenta/xvol-icm/17/17.html>.
- Antes, J. (1997). Structuring the process of airline scheduling. In P. Kischka, editor, *Operations Research Proceedings 1997*.
- Argüello, M., Bard, J., and Yu, G. (1997a). A GRASP for aircraft routing in response to groundings and delays. *Journal of Combinatorial Optimization*, **5**, 211–228.
- Argüello, M., Bard, J., and Yu, G. (1997b). Models and methods for managing airline irregular operations aircraft routing. In G. Yu, editor, *Operations Research in the Airline Industry*, pages 1–45. Kluwer Academic Publishers.
- Bard, J., Yu, G., and Argüello, M. (2001). Optimizing aircraft routings in response to groundings and delays. *IIE Transactions*, **33**, 931–947.
- Barnhart, C., Johnson, E., Nemhauser, G., and Vance, P. (1996). Exceptions crew scheduling. Technical report, The Logistics Institute, Georgia Institute of Technology.
- Barnhart, C., Johnson, E., Nemhauser, G., Savelsbergh, M., and Vance, P. (1998a). Branch-and-price: Column generation for solving huge integer programs. *Operations Research*, **46**, 316–329.
- Barnhart, C., Boland, N., Clarke, L., Johnson, E., Nemhauser, G., and Shenoi, R. (1998b). Flight string model for aircraft fleet and routing. *Transportation Science*, **32**, 208–220.
- Barnhart, C., Lu, F., and Shenoi, R. (1998c). Integrated airline schedule planning. In G. Yu, editor, *Operations Research in the Airline Industry*, pages 384–403. Kluwer Academic Publishers.
- Barnhart, C., Farahat, A., and Lohatepanont, M. (2002a). Airline fleet assignment: An enhanced revenue model. Technical report, Massachusetts Institute of Technology.
- Barnhart, C., Kniker, T., and Lohatepanont, M. (2002b). Itinerary-based airline fleet assignment. *Transportation Science*, **36**, 199–217.
- Barnhart, C., Cohn, A., Johnson, E., Klabjan, D., Nemhauser, G., and Vance, P. (2003). Airline crew scheduling. In R. W. Hall, editor, *Handbook of Transportation Science*. Kluwer Scientific Publishers.
- Bélanger, N., Desaulniers, G., Soumis, F., and Desrosiers, J. (2003). Periodic airline fleet assignment with time windows, spacing constraints, and time dependant revenues. Technical Report G-2003-41, GERAD.
- Benders, J. (1962). Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, **4**, 238–252.

- Berdy, P. (2002). Developing effective route networks. In D. Jenkins, editor, *Handbook of Airline Economics*. McGraw-Hill.
- Berge, M. (1994). Timetable optimization - formulation, solution approaches, and computational issues. In *AGIFORS proceedings*.
- Berge, M. and Hopperstad, C. (1993). Demand driven dispatch: A method for dynamic aircraft capacity assignment, models and algorithms. *Operations Research*, **41**, 153–168.
- Campbell, K., Durfee, R., and Hines, G. (1997). FedEx generates bid lines using simulated annealing. *Interfaces*, **27**, 1–16.
- Cao, J. and Kanafani, A. (1997a). Real-time decision support for integration of airline flight cancellations and delays part I: Mathematical formulations. *Transportation Planning and Technology*, **20**, 183–199.
- Cao, J. and Kanafani, A. (1997b). Real-time decision support for integration of airline flight cancellations and delays part II: Algorithms and computational experiments. *Transportation Planning and Technology*, **20**, 201–217.
- Carmen Systems. <http://www.carmen.se>.
- Chebalov, S. and Klabjan, D. (2002). Robust airline crew scheduling: Move-up crews. In *Proceedings of the 2002 NSF Design, Service, and Manufacturing Grantees Research Conference*.
- Clark, P. (2000). Dynamic fleet management. In G. Butler and M. Keller, editors, *Handbook of Airline Operations*, pages 273–285. McGraw-Hill.
- Clarke, L., Hane, C., Johnson, E., and Nemhauser, G. (1996). Maintenance and crew considerations in fleet assignment. *Transportation Science*, **30**, 249–260.
- Clarke, L., Johnson, E., Nemhauser, G., and Zhu, Z. (1997). The aircraft rotation problem. In R. E. Burkard, T. Ibaraki, and M. Queyranne, editors, *Annals of OR: Mathematics of Industrial Systems II*, pages 33–46. Baltzer Science Publishers.
- Clarke, M. D., Lettovský, L., and Smith, B. (2002). The development of the airline operations control center. In D. Jenkins, editor, *Handbook of Airline Economics*, pages 197–215. McGraw-Hill.
- Cohn, A. and Barnhart, C. (2002). Improving crew scheduling by incorporating key maintenance routing decisions. *Operations Research*. To appear.
- Cordeau, J., Stojković, G., Soumis, F., and Desrosiers, J. (2001). Benders decomposition for simultaneous aircraft routing and crew scheduling. *Transportation Science*, **35**, 375–388.
- Daskin, M. and Panayoyopoulos, N. (1989). A Lagrangian relaxation approach to assigning aircraft to routes in hub and spoke networks. *Transportation Science*, **23**, 91–99.
- Day, P. and Ryan, D. (1997). Flight attendant rostering for short-haul airline operations. *Operations Research*, **45**, 649–661.
- Desaulniers, G., Desrosiers, J., Dumas, Y., Solomon, M. M., and Soumis, F. (1997). Daily aircraft routing and scheduling. *Management Science*, **43**, 841–855.
- Desaulniers, G., Desrosiers, J., Ioachim, I., Solomon, M., and Soumis, F. (1998). A unified framework for deterministic time constrained vehicle routing and crew scheduling problems. In T. Crainic and G. Laporte, editors, *Fleet Management and Logistics*, pages 57–93. Kluwer Publishing Company.
- Desrosiers, J., Dumas, Y., Desrochers, M., Soumis, F., Sanso, B., and Trudeau, P. (1991). A breakthrough in airline crew scheduling. Technical Report G-91-11, GERAD.

- Desrosiers, J., Dumas, Y., Solomon, M., and Soumis, F. (1995). Time constrained routing and scheduling. In M. Ball, T. Magnanti, C. Monma, and G. Nemhauser, editors, *Handbook in Operations Research and Management Science, Network Routing*, pages 35–139. Elsevier Science Publishers.
- Ehrgott, M. and Ryan, D. (2001). Bicriteria robustness versus cost optimization in tour of duty planning at Air New Zealand. Technical report, Univeristy of Auckland.
- Erdmann, A., Nolte, A., Noltemeier, A., and Schrader, R. (1999). Modeling and solving the airline schedule generation problem. Available from <http://citeseer.nj.nec.com/erdmann99modeling.html>.
- Etschmaier, M. and Mathaisel, D. (1985). Airline scheduling: An overview. *Transportation Science*, **19**, 127–138.
- Feo, T. and Bard, J. (1989). Flight scheduling and maintenance base planning. *Management Science*, **35**, 1415–1432.
- Fisher, M. (1981). The Lagrangian relaxation method for solving integer programming problems. *Management Science*, **27**, 1–18.
- Fisher, M. (1985). An applications oriented guide to Lagrangian relaxation. *Interfaces*, **15**, 10–21.
- Galia, R. and Hjorring, C. (2003). Modeling of complex costs and rules in a crew pairing column generator. Technical Report CRTR-0304, Carmen Systems.
- Gamache, M. and Soumis, F. (1998). A method for optimally solving the rostering problem. In G. Yu, editor, *Operations Research in the Airline Industry*, pages 124–157. Kluwer Academic Publishers.
- Gamache, M., Soumis, F., Villeneuve, D., Desrosiers, J., and Gelinas, E. (1998). The preferential bidding system at Air Canada. *Transportation Science*, **32**, 246–255.
- Gamache, M., Soumis, F., Marquis, G., and Desrosiers, J. (1999). A column generation approach for large scale aircrew rostering problems. *Operations Research*, **47**, 247–262.
- Garvin, M. (2000). Service delivery system: A regional airline perspective. In G. Butler and M. Keller, editors, *Handbook of Airline Operations*, pages 419–428. McGraw-Hill.
- Gopalan, R. and Talluri, K. (1998). The aircraft maintenance routing problem. *Operations Research*, **46**, 260–271.
- Hane, C., Barnhart, C., Johnson, E., Marsten, R., Nemhauser, G., and Sigismondi, G. (1995). The fleet assignment problem: Solving a large-scale integer program. *Mathematical Programming*, **70**, 211–232.
- Jacobs, T., Johnson, E., and Smith, B. (1999). O&D FAM: Incorporating passenger flows into the fleeting process. In R. Darrow, editor, *Thirty-Ninth Annual AGIFORS Symposium*, New Orleans.
- Jarrah, A. and Diamond, J. (1997). The problem of generating crew bidlines. *Interfaces*, **27**, 49–64.
- Jarrah, A., Yu, G., Krishnamurthy, N., and Rakshit, A. (1993). A decision support framework for airline flight cancellations and delays. *Transportation Science*, **27**, 266–280.
- Kang, L. and Clarke, J. (2003). Degradable airline scheduling. Technical report, Massachusetts Institute of Technology.
- Karow, M. (2003). *Virtual Hubs: An Airline Schedule Recovery Concept and Model*. Master’s thesis, Massachusetts Institute of Technology.
- Kharraziha, H., Ozana, M., and Spjuth, S. (2003). Large scale crew rostering. Technical Report CRTR-0305, Carmen Systems.

- Klabjan, D., Johnson, E., Nemhauser, G., Gelman, E., and Ramaswamy, S. (2001a). Airline crew scheduling with regularity. *Transportation Science*, **35**, 359–374.
- Klabjan, D., Johnson, E., Nemhauser, G., Gelman, E., and Ramaswamy, S. (2001b). Solving large airline crew scheduling problems: Random pairing generation and strong branching. *Computational Optimization and Applications*, **20**, 73–91.
- Klabjan, D., Johnson, E., Nemhauser, G., Gelman, E., and Ramaswamy, S. (2002). Airline crew scheduling with time windows and plane count constraints. *Transportation Science*, **36**, 337–348.
- Kniker, T. (1998). *Itinerary-based airline fleet assignment*. Ph.D. thesis, Massachusetts Institute of Technology.
- Kohl, N. and Karisch, S. (2003). Airline crew rostering: Problem types, modeling, and optimization. *Annals of Operations Research*. To appear.
- Kwok, L. and Wu, L. (1996). Development of an expert system in cabin crew pattern generation. *International Journal of Expert Systems*, **9**, 445–464.
- Lavoie, S., Minoux, M., and Odier, E. (1988). A new approach for crew pairing problems by column generation with an application to air transportation. *European Journal of Operational Research*, **35**, 45–58.
- Lettovský, L. (1997). *Airline Operations Recovery: An Optimization Approach*. Ph.D. thesis, Georgia Institute of Technology.
- Lettovský, L., Johnson, E., and Smith, B. (1999). Schedule generation model. In R. Darrow, editor, *Thirty-Ninth Annual AGIFORS Symposium*, New Orleans.
- Lettovský, L., Johnson, E., and Nemhauser, G. (2000). Airline crew recovery. *Transportation Science*, **34**, 337–348.
- Listes, O. and Dekker, R. (2003). A scenario aggregation based approach for determining a robust airline fleet composition. Technical Report EI 2002-17, Rotterdam School of Economics.
- Lohatepanont, M. and Barnhart, C. (2002). Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment. Technical report, Massachusetts Institute of Technology.
- Løve, M., Sørensen, K., Larsen, J., and Clausen, J. (2002). Disruption management for an airline - rescheduling of aircraft. In S. Cagnoni, J. Gottlieb, E. Hart, M. Middendorf, and G. Raidl, editors, *Applications of Evolutionary Computing: EvoWorkshops 2002*, volume 2279 of *Lecture Notes in Computer Science*, pages 315–324. Springer-Verlag Heidelberg.
- Lučić, P. and Teodorović, D. (1999). Simulated annealing for the multi-objective aircrew rostering problem. *Transportation research part A*, **33**, 19–45.
- Makri, A. and Klabjan, D. (2003). Efficient column generation techniques for airline crew scheduling. *INFORMS Journal on Computing*. To appear.
- Marsten, R. (1994). Crew planning at Delta Airlines. XV Mathematical Programming Symposium, Ann Arbor. Presentation.
- Marsten, R., Subramanian, R., and Gibbons, L. (1996). Junior analyst extraordinare (JANE): Route development at Delta Air Lines. In *AGIFORS proceedings*.
- Martin, R. (1999). *Large Scale Linear and Integer Optimization : a Unified Approach*. Kluwer Academic Publishers.

- Mercier, A., Cordeau, J.-F., and Soumis, F. (2003). A computational study of Benders decomposition for the integrated aircraft routing and crew scheduling problem. Technical Report G-2003-48, GERAD.
- Minoux, M. (1984). Column generation techniques in combinatorial optimization: A new application to crew pairing problems. In *Proceedings XXIVth AGIFORS Symposium*.
- Minoux, M. (1986). *Mathematical Programming: Theory and Algorithms*. Wiley-Interscience.
- Paoletti, B., Cappelletti, S., and Lenner, C. (2000). Operations research models for the optimization of aircraft rotation and routing in the integrated resources management process. In G. Butler and M. Keller, editors, *Handbook of Airline Operations*, pages 285–308. McGraw-Hill.
- Rexing, B., Barnhart, C., Kniker, T., Jarrah, A., and Krishnamurthy, N. (2000). Airline fleet assignment with time windows. *Transportation Science*, **34**, 1–20.
- Rosenberger, J., Johnson, E., and Nemhauser, G. (2001). Rerouting aircraft for airline recovery. *Transportation Science*, **37**, 408–421.
- Rosenberger, J., Johnson, E., and Nemhauser, G. (2003). A robust fleet assignment model with hub isolation and short cycles. *Transportation Science*. To appear.
- Rushmeier, R. and Kontogiorgis, S. (1997). Advances in the optimization of airline fleet assignment. *Transportation Science*, **31**, 159–169.
- Ryan, D. and Foster, B. (1981). An integer programming approach to scheduling. In A. Wren, editor, *Computer Scheduling of Public Transport Urban Passenger Vehicle and Crew Scheduling*, pages 269–280. Elsevier Science B.V.
- Schaefer, A., Johnson, E., Kleywegt, A., and Nemhauser, G. (2000). Airline crew scheduling under uncertainty. Technical Report TLI-01-01, Georgia Institute of Technology.
- Sohoni, M. and Johnson, E. (2002a). Operational airline reserve crew planning. Technical report, The Logistics Institute, Georgia Institute of Technology.
- Sohoni, M. and Johnson, E. (2002b). An optimization approach to pilot recurrent training schedule. Technical report, The Logistics Institute, Georgia Institute of Technology.
- Sohoni, M., Johnson, E., and Bailey, G. (2003). Long range reserve crew manpower planning. Technical report, The Logistics Institute, Georgia Institute of Technology.
- Song, M., Wei, G., and Yu, G. (1998). A decision support framework for crew management during airline irregular operations. In G. Yu, editor, *Operations Research in the Airline Industry*, pages 260 – 286. Kluwer Academic Publishers.
- Sriram, C. and Haghani, A. (2003). An optimization model for aircraft maintenance scheduling and re-assignment. *Transportation Research Part A*, **37**, 29–48.
- Stojković, G., Soumis, M., and Desrosiers, J. (1998). The operational airline crew scheduling problem. *Transportation Science*, **32**, 232–245.
- Stojković, M. and Soumis, F. (2001). An optimization model for the simultaneous operational flight and pilot scheduling problem. *Management Science*, **47**, 1290–1305.
- Stojković, M. and Soumis, F. (2003). The operational flight and multi-crew scheduling problem. Technical Report G-2000-27, GERAD.
- Talluri, K. (1996). Swapping applications in a daily airline fleet assignment. *Transportation Science*, **30**.

- Teodorović, D. and Stojković, G. (1990). Model for operational daily airline scheduling. *Transportation Planning Technology*, **14**, 273–285.
- The Airline Group of the International Federation of Operational Research Societies (AGIFORS). <http://www.agifors.org>.
- Thengvall, B., Bard, J., and Yu, G. (2000). Balancing user preferences for aircraft schedule recovery during irregular operations. *IIE Transactions*, **32**, 181–193.
- Thengvall, B., Bard, J., and Yu, G. (2003). A bundle algorithm approach for the aircraft schedule recovery problem during hub closures. *Transportation Science*, **37**, 392–407.
- van Ryzin, G. and Talluri, K. (2002). Revenue management. In R. Hall, editor, *Handbook of Transportation Science*, pages 599–661. Kluwer Academic Publishers.
- Vance, P., Atamtürk, A., Barnhart, C., Gelman, E., Johnson, E., Krishna, A., Mahidhara, D., Nemhauser, G., and Rebello, R. (1997). A heuristic branch-and-price approach for the airline crew pairing problem. Technical Report LEC-97-06, Georgia Institute of Technology.
- Wei, G. and Yu, G. (1997). Optimization model and algorithm for crew management during airline irregular operations. *Journal of Combinatorial Optimization*, **1**, 305–321.
- Yan, S. and Lin, C. (1997). Airline scheduling for the temporary closure of airports. *Transportation Science*, **31**, 72–82.
- Yan, S. and Tseng, C. (2002). A passenger demand model for airline flight scheduling and fleet routing. *Computers and Operations Research*, **29**, 1559–1581.
- Yan, S. and Wang, C. (2001). The planning of aircraft routes and flight frequencies in an airline network operations. *Journal of Advanced Transportation*, **35**, 33–46.
- Yan, S. and Young, H. (1996). A decision support framework for multi-fleet routing and multi-stop flight scheduling. *Transportation Research Part A*, **30**, 379–398.
- Yen, J. and Birge, J. (2000). A stochastic programming approach to the airline crew scheduling problem. Technical report, University of Washington.