

Commissioned Paper

Airline Schedule Planning: Accomplishments and Opportunities

Cynthia Barnhart

Operations Research Center, and the Center for Transportation and Logistics, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02134, cbarnhart@mit.edu

Amy Cohn

Industrial and Operations Engineering, University of Michigan, Ann Arbor, Michigan 48109, amycohn@umich.edu

Plagued by high labor costs, low profitability margins, airspace and airport congestion, high capital and operating costs, security and safety concerns, and complex and large-scale management and operations decisions, the airline industry has armed its planners with sophisticated optimization tools to improve decision making and increase airline profits. In this paper, we describe optimization approaches for airline schedule planning, demonstrating how optimization can facilitate the management of a diverse and finite set of expensive, highly constrained resources. We focus on the art and science of modeling and solving these problems, providing illustrative examples of the associated impacts and challenges, and highlighting effective techniques that might be applicable to problems arising in other industries.

Key words: airline scheduling; network design; large-scale optimization

History: Received: November 7, 2002; days with authors: 115; revisions: 2; average review cycle time: 40 days; Senior Editor: Leroy Schwarz.

1. Introduction: Optimization and the Airline Industry

Driven by competitive pressures, low profit margins, and complex operations involving expensive assets and crews, the airline industry has been a leader in the development and application of operations research methods. As a case in point, in the late 1970s, when the U.S. domestic airline industry was deregulated, existing U.S. airlines were threatened by intense competition from low-cost, new entrants such as People's Express and Southwest Airlines. The existing airlines faced the grim choice of matching the fares charged by the new entrants and losing money, or not matching the new entrant fares and losing market share. Either choice produced the same end result—The existing airlines faced elimination. Instead, a number of airlines turned to operations research for an alternative. The answer was revenue management. By creating multiple fares and allocating seat inventory among the fare classes, exist-

ing airlines could match the low-cost carrier fare for a portion of their seat inventory and charge higher fares for high-value customers (typically business travelers). Because the new entrants did not have the operations research and information technology expertise needed to design and implement revenue management systems, existing carriers were able to survive deregulation and the challenge of low-cost entrants.

Revenue management represents only one of many operations research success stories for airlines. Another is airline schedule planning. Schedule planning involves designing future airline schedules to maximize airline profitability. In designing schedules, the airline answers questions of which origin-destination markets to serve, with what frequency, and over which hubs. Moreover, the airline determines the departure time and aircraft type for each of these flights. With thousands of flights per day; hundreds of aircraft of differing types; hundreds of airports and multiple hubs; constraints related to

gates, airport slots, air traffic control, curfews, maintenance, and crews; and complex issues involving competition, demand forecasting, pricing, and revenue management, schedule planning poses daunting challenges for the schedule planner. Applying sophisticated forecasting and optimization tools is critical. American Airlines claims that their schedule-planning system generates over \$500 million in incremental profits annually. This impact is deeply significant when you consider that without revenue management and schedule-planning systems in the period 1990–2000, American Airlines would have been profitable in only one year, an all-time record year for profitability (Cook 2000). Other testimonials to the impact of optimization of the airline industry can be found in Yu et al. (2003), Butchers et al. (2001), Wiper et al. (1994), Smith et al. (1992), and Patty et al. (1991).

In this paper, we describe the state-of-the art and practice in airline schedule planning optimization. While this focused topic does not provide a comprehensive view of the operations research challenges faced by airlines, it does provide insights into the design of effective optimization approaches for airline planning, and their impacts on the industry. Our goal is to illustrate effective techniques for modeling and solving airline scheduling problems that might be applicable to other industries.

In §2, we describe the schedule-planning problems airlines face, overview various modeling approaches, and highlight the impacts and challenges associated with these efforts. In §3, we focus on the blend of art and science needed to solve these large-scale, complex schedule-planning models, again highlighting effective approaches, impacts, and challenges. Finally, in §4, we describe future research opportunities and challenges.

2. Scheduling Problems

Before an airline begins operation, it has to resolve a number of questions, including:

(1) *Schedule Design*. Which markets to serve and with what frequency, and how to schedule flights to meet these frequencies.

(2) *Fleet Assignment*. What size aircraft to assign to each flight.

(3) *Aircraft Maintenance Routing*. How to route aircraft to ensure the satisfaction of maintenance requirements.

(4) *Crew Scheduling*. Which crews to assign to each flight so that crew costs are minimized.

Collectively, these problems are characterized by numerous and well-defined constraints, a large number of interrelated decisions, and nonlinear cost structures. To illustrate this, consider that (1) revenues associated with assigning an aircraft to one flight leg often cannot be computed without knowing the types of aircraft assigned to *several other* flight legs in the network (because passenger itineraries often include multiple flights) and (2) for a particular flight leg, the cost of the assigned crew depends on which *other* flight legs are assigned to that crew (and thus the amount of time spent on the ground between flights). To model the interdependencies and nonlinearities of the above problems, very large-scale linear programs with integer variables have been developed, typically containing *at least* thousands of constraints and billions of decision variables. The combination of problem complexity and size often renders decision makers ineffective in producing cost-effective solutions, despite their wealth of intuition and experience. The airlines have thus turned to optimization techniques.

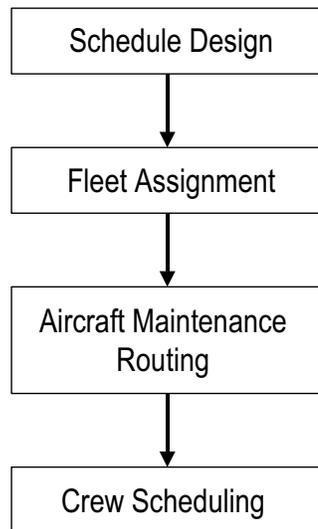
To date it has been impossible to construct and solve a single optimization model that alone can answer this myriad of questions. One key difficulty stems from model size: A single model is simply too large, containing many billions of decision variables and constraints. Airlines have instead employed a decomposition approach, breaking up the large problem into a set of smaller problems and then solving them sequentially. Even these smaller *core* problems, resulting from the sequential solution approach, present significant modeling and solution challenges, as discussed below.

2.1. The Core Problems

Schedule planning at airlines is typically broken into a set of four core problems, each solved sequentially. These problems, depicted in Figure 1, are *schedule design*, *fleet assignment*, *aircraft maintenance routing*, and *crew scheduling*.

2.1.1. Schedule Design. The flight schedule, containing the flight legs to be flown and the departure time of each flight leg, is the single most important product of an airline. It largely defines the market share an airline will capture, and hence is a key

Figure 1 Aircraft and Crew Schedule Planning



determinant of airline profitability. Designing a flight schedule to maximize profitability is extraordinarily complex, with essentially all elements of the airline (and competing airlines as well) linked to the flight schedule design decisions. One leading industry executive (Cook 2000) remarked, “The network planners are usually the smartest people in an airline.”

Impacts and Challenges. Operations researchers, armed with optimization models and methods, have not fully met the challenges posed by the size and complexity of airline schedule design, for reasons including:

(1) *Complexity and Problem Size.* Decisions involving schedule design are integrally tied to the assignment of aircraft and crews to the flight legs. Incorporating crew, fleet, and schedule design decisions into a single model, however, involves numerous flight-leg options and complex constraints and interdependencies. Moreover, such a model would be intractable, containing many billions of constraints and variables.

(2) *Data Availability and Accuracy.* Critical inputs for designing the flight schedule include the airline’s expected *unconstrained market demands* and *average fares*, two particularly challenging forms of data.

- The airline’s unconstrained market demand is defined as the maximum demand the airline will experience for an origin-destination market, independent of the capacity provided in that market. Unconstrained demand data are critical to developing a

good flight schedule, but extremely difficult to estimate accurately. The first issue is that unconstrained demands cannot be directly observed. The second is that the demand an airline will experience in a given market is a function of the airline’s schedule in that market, with greater frequency of service typically leading to greater market share. This relationship creates a chicken-and-egg effect, with demand for airline service a function of the airline schedule and the optimal airline schedule a function of market demand. Moreover, an airline’s demand is influenced by the schedule of competing airlines, again inducing a chicken-and-egg effect, with competing airlines adjusting their schedules as an airline makes changes to its own schedule. Air carrier demand research appears in Prousaloglou and Koppelman (1995, 1999).

- Average fares are another core element of the input data for schedule design. (Because an airline offers an enormous number of fares, models are typically based on *average fares*.) Average fares are affected by complex pricing and revenue management systems, the order in which customers request tickets for a flight, and the demand and capacity in a market—which are themselves dependent on the flight schedule. Further complicating the estimation of fares are the competitive pressures in the industry that force airlines to adjust their fare structures dynamically, reacting to pricing changes made by their competition. Industry executives lament that an airline’s pricing strategy can be only as good as its worst competitor.

These compounding difficulties have contributed to the current, typical airline practice of designing flight schedules manually, with limited optimization. However, by limiting the complexity and scope of models, recent progress has been made in applying optimization to schedule design problems. Researchers have focused on determining incremental changes to flight schedules, producing a new schedule by applying a limited number of changes to the existing schedule. Generally, the approach is to cast these problems as network design models, with the objective of maximizing the incremental profits of the new schedule.

In their incremental optimization approach, Lohatepanont and Barnhart (2001) select flight legs to include in the flight schedule and simultaneously

optimize aircraft assignments to these legs. Using their approach—an extension of the fleet assignment models and algorithms in the following sections—they demonstrate potential improvements in aircraft utilization and significant increases in revenue, with an estimated impact exceeding \$200 million at one major airline. Other recent research efforts to optimize airline schedule design are described in Armacost et al. (2002) and Erdmann et al. (1999).

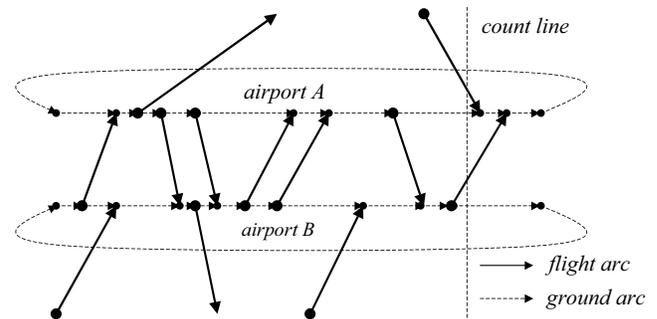
2.1.2. Fleet Assignment. Once the flight schedule has been specified, the next step in schedule planning is *fleet assignment*. The fleet assignment problem is to assign a particular *fleet* (equipment) type to each flight leg so as to minimize *operating* and *spill* costs. Operating costs vary by aircraft type and represent the cost of flying a flight leg with an aircraft of that type. Spill costs represent the revenue lost when passenger demand for a flight leg exceeds the assigned aircraft's seating capacity, resulting in *spilled demand* (that is, passengers turned away by the airline) and lost potential revenue.

At first glance, the approach to determining optimal fleet assignments appears straightforward—simply match as closely as possible flight demand with assigned aircraft seating capacity. The optimal solution is complicated, however, by the interactive nature of the flight network. Specifically, assigning an aircraft to one leg of the network impacts the feasible assignments of aircraft to other flight legs. This occurs both because the airline has a limited number of aircraft of each type and because assignments of aircraft to flight legs must be balanced, with the number of aircraft of a particular type arriving at a location equal to the number departing that location.

Modeling Approach. To capture these considerations, fleet assignment models have historically been cast as multicommodity network flow problems with side constraints (Hane et al. 1995). The underlying network (see Figure 2 for an illustration) is a *time-line network*, with

- (1) Nodes corresponding to the times and locations of flight departures and arrivals.
- (2) Two types of arcs, *flight arcs* and *ground arcs*. Flight arcs represent the schedule's flight legs, with arrival times adjusted to include the minimum amount of time aircraft must remain on the ground.

Figure 2 A Time-Line Network Involving Two Airports



This time is needed for tasks such as disembarking and embarking passengers, unloading and loading baggage, and refueling. Ground arcs represent aircraft on the ground between flights.

Each aircraft type is a commodity, and the objective is to flow each commodity through the network feasibly and with minimum cost, such that each flight leg is assigned to exactly one aircraft type. For each feasible assignment of a fleet type to a flight leg, the decision of whether or not to make that assignment is represented by a binary variable.

Impacts and Challenges. Fleet assignment models are widely used by the industry in making their fleet decisions and are credited with contributing to significant cost savings, measuring in the millions of dollars annually. For example, Rushmeier and Kontogiorgis (1997) state that at USAir, an audit based on *actual* results indicates savings of at least \$15 million annually. Using fleet assignment models, Wiper et al. (1994) report annual savings of \$100 million at Delta Airlines, and Abara (1989) reports a 1.4% improvement in operating margins at American Airlines. In addition to providing economic benefits, research on fleet assignment problems, such as that presented in Hane et al. (1995), led to advanced techniques for solving large-scale linear programs. Many of these techniques have been generalized and incorporated into commercial solvers, allowing more efficient solution of large-scale optimization problems.

These impressive results notwithstanding, there remain several critical challenges in fleet assign-

ment. Many of these challenges stem from modeling assumptions that include:

(1) Many fleet assignment models assume that the flight schedules repeat daily, even though most airlines operate different schedules on the weekend.

(2) Most fleet assignment models assume flight-leg demand is known and does not vary by day of week, but historical data clearly show that day-to-day demand variations are present.

(3) Flying times and ground times are typically assumed to be deterministic in fleet assignment models; however, historical data show large variations in flight and ground times, caused by factors including congestion on the ground and in the air, weather conditions, and new security practices.

(4) Most fleet assignment models assume that the number of spilled passengers, and their associated spill costs, can be computed per *flight leg*. In fact, passenger demand, spill, and the revenue associated with each passenger is *itinerary* specific, not *flight-leg* specific. Hence, obtaining unconstrained demand and average fare for each *flight leg* in order to compute spill costs poses challenges, as described in the next section.

Estimates of Spill Costs. Greatly limiting the accuracy of many fleet assignment models is the fact that spill costs must be estimated for each *flight leg*, and not for each *itinerary*. To achieve this, itinerary fares must be allocated to flight legs so that the total resulting leg-based spill costs equal total spill costs over all itineraries. One allocation rule satisfying this criterion is to allocate an itinerary’s fare to its flight legs such that their sum equals the total fare of the itinerary. While necessary, this rule is not sufficient to ensure that the spill costs of a particular fleet assignment will be correctly computed. The issue stems from the multileg nature of some itineraries.

Consider, for example, an airline that operates a network consisting of flight legs X–Y and Y–Z, serving two nonstop origin-destination markets X–Y and Y–Z, and one connecting market X–Z, with flight legs X–Y and Y–Z. Table 1 shows unconstrained market demands and average fares for these markets. Assume that the airline operates two aircraft types, one with 100 seats and the other with 200 seats. Further assume that the scheme to allocate itinerary fares

Table 1 Demand and Fare Data

Market	Itinerary (Sequence of Flight Legs)	Number of Passengers	Average Fare (\$)
X–Y	1	75	200
Y–Z	2	150	225
X–Z	1–2	75	300

to flight legs results in the following spill costs:

- (1) \$200 for the X–Y market on the X–Y flight leg.
- (2) \$225 for the Y–Z market on the Y–Z flight leg.
- (3) \$150 for the X–Z market on the X–Y flight leg.
- (4) \$150 for the X–Z market on the Y–Z flight leg.

Consider the assignment of the 100-seat aircraft to both flight leg 1 (X–Y) and flight leg 2 (Y–Z). Given these allocated costs, the optimal spill and its associated cost for each flight leg can be characterized as:

Flight Leg 1. Spill 50 X–Z passengers, for a total associated spill cost of $50 * \$150$.

Flight Leg 2. Spill 50 Y–Z passengers and 75 X–Z passengers on flight leg 2, for a total spill cost of $50 * 225 + 75 * 150$.

This simple example illustrates the inconsistencies that result with a leg-based approach to fleet assignment; namely, spill costs are computed using the assumption that spill on each leg can be optimized without regard to other flight legs in the network. In this example, all 75 passengers from X–Z are spilled from flight leg 2, but only 50 are spilled from flight leg 1. Clearly, this results in an underestimation of true spill—passengers do not travel on only a portion of their desired itineraries.

Another issue of this leg-based approach is that spilling passengers in an optimal manner on each *flight leg* based on the allocated spill costs does not result in an optimal spill solution. Again, with 100-seat aircraft assigned to both flight legs 1 and 2, the true optimal spill is 75 Y–Z passengers and 50 X–Z passengers. It is cost effective to displace from flight leg 2 a Y–Z passenger with allocated spill cost of \$225 and replace that passenger with an X–Z passenger with allocated spill cost of \$150 because the X–Z passenger also has allocated cost of \$150 on flight leg 1. A leg-based approach cannot recognize such a solution as optimal, because it does not capture these network interdependencies. In fact, Barnhart et al.

(2002b) show that there is no single scheme for allocating itinerary fares to flight legs that can always guarantee an optimal fleet assignment.

To overcome the inaccuracies of leg-based models, more recent modeling efforts have introduced itinerary- or origin-destination-based fleet assignment models (Barnhart et al. 2002b, Erdmann et al. 1999, Jacobs et al. 1999). In Barnhart et al. (2002b), the fleet assignment model described above is expanded to include path variables that represent the assignment of passengers to itineraries, obviating the need to allocate itinerary fares to flight legs. The passenger flow variables are linked to the aircraft assignment variables through the capacity constraints, ensuring that the number of available seats is not exceeded by the number of assigned passengers. In a case study involving a major U.S. airline, they report that fleet assignments using network-based models with passenger itineraries improve upon leg-based fleet-assigning solutions significantly, with estimated savings ranging from 30 million to over 100 million dollars annually.

2.1.3. Aircraft Maintenance Routing. Once the fleet assignment decisions have been made, the flight network decomposes into subnetworks, each one associated only with the aircraft of a particular fleet type. The aircraft-maintenance-routing and crew-scheduling steps of the schedule-planning process, therefore, require the solution of much smaller problems, one for each aircraft type.

Given the subnetwork of flight legs assigned to a particular aircraft type, the assignment of *individual* aircraft to these flight legs occurs in the *aircraft-maintenance-routing* step. The goal of aircraft maintenance routing is to determine *routings*, or *rotations*, for each aircraft in a fleet. A routing is a sequence of flight legs, with the destination of one flight leg the same as the origin of the next leg in the sequence. A rotation is a routing which starts and ends at the same location. Each aircraft's rotation is mandated to visit, at regular intervals, airports that serve as maintenance stations for that aircraft type. This mandate is motivated by regulatory restrictions requiring aircraft to undergo maintenance checks between blocks of flying time, not to exceed a specified limit. Detailed descriptions of

the maintenance-routing problem are contained in Gopalan and Talluri (1998) and Clarke et al. (1996b).

Modeling Approach. The aircraft-maintenance-routing problem can be modeled as a *network circulation problem* with side constraints. The decision variables correspond to sequences (*strings*) of flight legs, with the sequence (1) beginning and ending at a maintenance station and (2) satisfying the rules governing the maximum time between maintenance. If a string is included in the solution, a single aircraft flies each flight in the sequence and then undergoes maintenance. Side constraints, beyond the flow circulation constraints ensuring aircraft balance, include *cover constraints* and *count constraints*. Cover constraints ensure that each flight leg is contained in exactly one selected string, and count constraints limit the number of assigned aircraft to the number available. Additional details are provided in Barnhart et al. (1998a).

The primary objective of the aircraft-maintenance-routing problem is to find a *feasible* solution, one ensuring sufficient maintenance opportunities for each aircraft. Historically, another objective has been to determine routings that maximize *through revenue*. Through revenue is the incremental fare airlines charge passengers for the convenience of staying on the same aircraft between flights. When computer reservation systems (CRSs) were the primary distribution channel, prior to the use of the Internet as a way to purchase tickets, *through itineraries* (with passengers staying on the same aircraft between flights) were given priority in CRS displays over connecting itineraries (with passengers changing aircraft between flights). Not surprisingly, the order in which an itinerary appeared on a CRS display was correlated with the number of sales of that itinerary, with itineraries appearing first having the highest booking levels. Using the Internet, airline passengers are increasingly aware of available itineraries and their associated fares. Arguably, given the price sensitivity of many passengers, this diminishes the relevance of through revenues. Moreover, because of disruptions, aircraft rotations are often altered during operations and passengers who were booked on through itineraries are in fact required to change planes to connect. These factors suggest that in modeling today's aircraft-maintenance-routing problems, the objective

might best be cast as one only of achieving a feasible maintenance routing.

Impacts and Challenges. Barnhart et al. (1998) note that aircraft-routing problems of major U.S. airlines, containing up to 200 flight legs and over 400 million string variables, can often be solved on workstation-class computers, typically requiring less than 10 hours of computation time. This suggests that feasible solutions can be found, *when they exist*. Feasibility is not guaranteed, however, given the sequential approach of first solving the fleet assignment problem and then solving the resulting aircraft-routing problems. Barnhart et al. (1998) found the likelihood of a fleet assignment solution yielding maintenance-feasible aircraft routings to be a function of the structure of the flight network.

Sequential Solution and Network Structure. In *hub-and-spoke networks*, the effects of a sequential approach are mitigated by the existence of *banks*, that is, times at which many aircraft are on the ground at a hub airport to allow passengers to connect between flights. With many aircraft of a given type at the same location at roughly the same time, banks provide aircraft with many possible routings. This is in contrast with *point-to-point networks*, especially international networks with low frequencies. Separating the fleet assignment and aircraft-routing decisions in this latter case can often result in fleetings without maintenance-feasible aircraft routings.

The likelihood that the fleet assignment solution yields maintenance-infeasible aircraft routings is reduced by incorporating approximate maintenance constraints, called *pseudomaintenance constraints*, into the fleet assignment model. Ensuring that sufficient numbers of aircraft of each type are located at maintenance stations nightly, these constraints are straightforward to include in fleet assignment models. However, they cannot guarantee that each *individual* aircraft will have sufficient maintenance opportunities. Instead, pseudomaintenance constraints enforce *aggregate* requirements on the number of aircraft that must be maintained daily. Indeed, one aircraft might be maintained every night while another is never maintained.

To guarantee the satisfaction of *individual* aircraft maintenance requirements, the fleet assignment

model must be altered dramatically, with decision variables based on strings (as in aircraft-routing models). Expanding the model in this manner, effectively integrating the fleet assignment and aircraft-routing models, increases the number of variables dramatically and introduces tractability issues (Barnhart et al. 1998c).

For hub-and-spoke airlines, this added complexity is rarely warranted. Including pseudomaintenance constraints in the fleet assignment models is almost always sufficient to ensure a feasible maintenance routing. For point-to-point airlines, however, the altered model is often necessary to ensure that the resulting solution is maintenance feasible. Additional detail about the integrated fleet assignment and aircraft routing problem is provided in §2.2.

Euler-Tour Requirement. An interesting twist often applied to aircraft routings in U.S. domestic networks is the *Euler-tour requirement*. In order to ensure equal wear and tear on their aircraft, several airlines require that *all* aircraft of a particular type fly all flight legs assigned to that aircraft type. This requirement can be formulated as the problem of finding one maintenance-feasible rotation containing each flight leg exactly once (an Euler tour) and requiring no more than the number of aircraft of that fleet type to fly all flight legs daily. Again, the difficulty of satisfying this particular constraint is a function of the network structure of the flight schedule.

In hub-and-spoke networks, this restriction is typically easily satisfied through a sequential solution approach, first solving fleet assignment with pseudomaintenance constraints and then solving aircraft routing with Euler-tour constraints. In contrast, to ensure feasible solutions in point-to-point networks, it is often necessary to build integrated models capturing the fleet assignment, maintenance routing, and Euler constraints simultaneously. For a more detailed discussion of this topic, the reader is referred to Barnhart et al. (1998c).

2.1.4. Crew Scheduling. Airlines typically partition their crew-scheduling problem into two categories. *Cabin crews* (flight attendants) are responsible for the comfort and safety of the passengers, and *cockpit crews* (pilots) are responsible for operating the aircraft and executing the flight leg. Most optimization

research on airline crew scheduling has focused on the pilot problem, primarily because cockpit crews are paid substantially more than cabin crews. The cockpit crew problem is also smaller and simpler in structure than the cabin crew problem, for reasons including:

(1) Cockpit crews, unlike cabin crews, stay together throughout the workday, allowing crews, instead of individuals, to be modeled.

(2) For each flight leg, the size and composition of the cockpit crew is known a priori, based upon the defined requirements for each aircraft. The cabin crew size, however, varies as a function of the number of passengers on the flight leg.

(3) Cockpit crews have fewer options: They can be assigned only to those flight legs with assigned aircraft types that they are qualified to fly. Cabin crews have much more flexibility, with only limited restrictions on the type of aircraft to which they are assigned.

For both cockpit or cabin crews, the crew-scheduling problem itself is usually solved sequentially. First, a set of minimum-cost work schedules (*pairings*) typically spanning one to five days is determined. Feasible pairings are constrained by a host of work rules imposed by regulatory agencies and collective bargaining agreements. Rules include restrictions on the maximum number of work hours in a day, the minimum number of hours of rest between work periods, the maximum time away from home, etc. These restrictions notwithstanding, the number of possible pairings often numbers in the billions for major U.S. carriers. Further complicating the problem is the cost of a pairing, usually represented as a non-linear function of flying time, total elapsed work time, and total time away from home.

The second step in solving crew-scheduling problems is to assemble pairings into monthly work schedules, called *bidlines* or *rosters*, and to assign them to individual crew members. With rostering, a practice commonly used outside the United States, schedules are constructed for and assigned to specific individuals, taking into consideration their particular needs or requests. This is in contrast to bidline generation—a practice more commonly used in the United States. With bidline generation, schedules are generated but not for specific individuals.

The cost-minimizing subset of schedules is selected, and individual employees reveal their relative preferences for these schedules through a *bidding* process. The airline then makes the specific assignment of schedule to employees based on individual priority rankings—rankings that are often related to seniority. More details on the bidline and rostering problems can be found in Kohl and Karisch (2002), Cappanera and Gallo (2001), Caprara et al. (1998), Christou et al. (1999), Dawid et al. (2001), Day and Ryan (1997), and Gamache et al. (1998).

In addition to the complexity already described, crew-scheduling models must also incorporate a host of other constraints, including crewbase requirements, work balance concerns, etc. Moreover, depending on whether the flight schedule is domestic or international, different modeling and solution techniques are required. Detailed descriptions of the airline crew-scheduling problem are included in the survey papers by Barnhart et al. (2003), Clarke and Smith (2000), and Desaulniers et al. (1998).

Modeling Approach. The crew-pairing, bidline, and rostering problems can all be cast as *set-partitioning problems*. While simple in form, the set-partitioning model is a powerful one in this context. To illustrate, we focus on the crew-pairing problem. Construct one binary decision variable y_p for each pairing $p \in P$ (where P is the set of all feasible pairings), and define δ_{fp} to be an *indicator variable*, with value one if pairing p contains flight $f \in F$ (where F is the set of all flight legs) and zero otherwise. The crew-pairing problem is then formulated as

$$\min \sum_{p \in P} c_p y_p, \quad (1)$$

$$\text{subject to } \sum_{p \in P} \delta_{fp} y_p = 1 \quad \forall f \in F \quad (2)$$

$$y_p \in \{0, 1\} \quad \forall p \in P. \quad (3)$$

The objective (1) minimizes the cost of the chosen set of pairings. The cover constraints (2) and integrality constraints (3) require that for each flight f , exactly one pairing is chosen containing that flight.

By defining variables to correspond to *sequences of flight legs*, only *feasible* pairings are considered and explicit representation of complicated work rules

is unnecessary. Moreover, nonlinear costs associated with a day of work are captured by a *single quantity* representing pairing cost. The result is a crew-pairing model with simple form, containing a linear objective function and binary decision variables. With its advantages, however, comes one major disadvantage—namely, the huge number of possible crew pairings (often, numbering in the billions or more). Hub-and-spoke network structures, in particular, exacerbate this issue. With many aircraft on the ground at the same time and location, a huge number of crew connections, and hence pairings, is possible. In §3, we present some of the specialized algorithmic approaches used to overcome this obstacle.

Impacts and Challenges. Airline crew scheduling has garnered the most attention of optimizers, with research spanning decades. Crew-scheduling problems, with their numerous, complex rules and well-defined costs, are particularly amenable to optimization. Because there are so many possible decisions, it can be difficult to find feasible—let alone close to optimal—solutions manually. Moreover, crew costs represent the airlines' second-highest operating cost after fuel, so even slight improvements in the utilization of crews can translate into significant savings.

For these reasons, most major airlines use optimization tools to generate part or all of their crew schedules. As is the case with aircraft-scheduling tools, crew-scheduling optimization systems are either developed in-house by the airline's operations research group (often in collaboration with researchers from academia), or purchased from companies with expertise in airline optimization, including other airlines. Optimization has had a significant economic impact on crew scheduling. A decade ago solutions to crew-pairing problems were typically 10%–15% above the lower bound of flying cost, compared to today's solutions which typically are within at most 1%–2% of the lower bound. This improvement in solution quality translates to savings on the order of \$50 million annually for a large airline (Barnhart et al. 2003).

In addition to these economic benefits, crew-pairing optimization can be a useful negotiation tool, helping airlines to quantify the impacts of proposed changes in cost structures, benefits, and work rules proposed by their labor unions. The economic impact of each

proposed change can be evaluated precisely by changing selected inputs to the crew-pairing model, and re-running the optimization procedure.

While significant, the effects of optimization on crew scheduling are nonetheless limited by the sequential nature of the solution process. Once the flight schedule, the fleet assignments, and the aircraft-routing decisions have been made, the range of crew possibilities is significantly narrowed. Although researchers have been unsuccessful to date in modeling and solving the fully integrated schedule-planning problem, they are developing models that integrate pieces of the overall problem, as described in the next section.

Another issue associated with optimized crew schedules (and, in fact, with optimized fleet assignments and maintenance routings as well) is that the true impact of optimization is not exactly known. One reason for this is that the optimized solutions are rarely executed as planned. Crew sickness, vacation, and training periods, as well as mechanical failures and adverse weather conditions, result in necessary changes, often leading to significantly increased costs. Moreover, a finely tuned, optimized solution achieves increased utilization through the removal of slack, providing crews with less time to connect between flights and aircraft with less time on the ground between flying. Less slack time, although economical in theory, can translate in practice into less robustness and increased costs. In §4, we further discuss this important issue.

2.2. Integrating Core Models

After decades of research, *opportunities* to improve core problem models still exist, but *improvements*, especially significant ones, are increasingly difficult to achieve. Some researchers have thus turned their attention to modeling integrated core problems, in an attempt to ameliorate some of the drawbacks of a sequential solution approach. Because a tractable model integrating *all* decisions is not currently attainable, researchers have adopted alternative strategies. These strategies achieve *partial integration* of the core problems either by:

- (1) Merging two core problem models to create an integrated model that fully captures the two problems or

(2) Enhancing a core problem model by incorporating into it some of the key elements of another core problem model.

EXAMPLE 1: THE INTEGRATED FLEET ASSIGNMENT AND AIRCRAFT-MAINTENANCE-ROUTING PROBLEM. Barnhart et al. (1998c) ensure that fleet assignment decisions guarantee maintenance feasibility by merging the fleet assignment and aircraft-routing models. Their model contains one core aircraft-routing model for each aircraft type, along with the fleet assignment constraints requiring that each flight leg be assigned to exactly one aircraft type. They show that for a set of problems based on an international flight schedule containing more than 1,000 flight legs and nine fleet types, five out of nine had no maintenance-feasible solution when the fleet assignment and routing problems were solved sequentially. Their integrated model, however, produced near-optimal fleet assignment and feasible routing solutions in less than six hours, using a workstation-class computer.

EXAMPLE 2: THE INTEGRATED AIRCRAFT-MAINTENANCE-ROUTING AND CREW-PAIRING PROBLEM. In the sequential solution approach, the aircraft-routing solution imposes limitations on the possible crew-scheduling opportunities, resulting in potentially significant increases in crew costs. To understand the connection between the two problems, it is important to understand that crews and aircraft are independent. A crew connecting between two flights might stay on the same aircraft, or might transfer between two different aircraft, depending on the aircraft routings. Transferring between different aircraft requires more time (thereby introducing inefficiencies) than staying with the same aircraft. Hence, the set of feasible crew solutions changes as aircraft rotations change. Specifically, a crew cannot be assigned to a *short connect* (a flight connection with less than the minimum required crew connection time) unless both of the flights in the connection are assigned to the same aircraft. This highlights one of the key issues of sequential solution: The aircraft-routing solution is optimized without consideration of associated crew costs, yet crew costs are the airline's second-largest operating costs.

Klabjan et al. (2002) address this by solving the crew-pairing problem *before* the maintenance-routing

problem, and provide computational results demonstrating the degree to which choosing the “wrong” maintenance-routing solution can impact the quality of the crew solution. Although this approach has the benefit of focusing on minimizing crew costs, it does not ensure that a maintenance-feasible solution will be found. To guarantee compatible solutions, Cordeau et al. (2000) directly integrate the basic maintenance-routing and crew-scheduling models. As an alternative approach, Cohn and Barnhart (2003) achieve improved crew-scheduling solutions by integrating *key elements* of the maintenance-routing and crew-scheduling problems. In their work, the crew-pairing problem and maintenance-routing problems are merged, but the aircraft-routing model is significantly modified to contain only decisions that (1) are relevant to the crew-pairing problem, and (2) assure a feasible maintenance-routing solution. We provide added discussion of this problem in §3.

EXAMPLE 3: ENHANCED FLEET ASSIGNMENT TO CAPTURE THE IMPACT ON CREW COSTS. In a sequential solution approach, crew-pairing problems decompose by aircraft type, allowing one problem to be solved independently for each type. Merging the crew-pairing and fleet assignment models, however, results in an integrated and *intractable* model that includes *all* crew pairings for *all* fleet types. Nonetheless, airline schedule planners have keen interest in improving upon the sequential solution approach, with its costly impact on crew scheduling.

Clarke et al. (1996a) and Barnhart et al. (1998a) have taken steps towards integrating fleet assignment and crew scheduling by modifying fleet assignment models to account for some of the downstream effects on crews. Clarke et al. (1996a) include *bonuses* for fleet assignments that allow crews to stay with the same aircraft. This reduces the occurrence of extended periods of crew rest between flights (the fleet assignment objective indirectly works to achieve solutions with high aircraft utilization), thereby eliminating unnecessary crew inactivity and costs.

In Barnhart et al. (1998a) crew costs are more explicitly incorporated into the fleet assignment model by adding decision variables representing *partial* crew pairings, with *approximate* costs. Solving their enhanced model using data from a long-haul airline, they achieve a reduction in crew costs of over 7% with

a corresponding increase in fleet costs of just over 1%, for a total decrease in combined crew and fleet costs of over 3%.

EXAMPLE 4: ENHANCED FLEET ASSIGNMENT TO INCLUDE SCHEDULE DESIGN DECISIONS. Because the core schedule design problem has not been successfully modeled in all its complexity, researchers have focused on determining incremental changes to flight schedules to allow a better matching of flight schedules and fleetings. In one such approach, proposed by Rexing et al. (2000), the fleet assignment model is modified to allow small retimings of flight-leg departures (on the order of 5 to 20 minutes). This retiming can allow additional aircraft assignments. To understand how, consider a spoke station with only two flights, one inbound and the other outbound, throughout the day. If the scheduled arrival time of the inbound flight and the scheduled departure time of the outbound flight are separated by only 25 minutes, these two legs might require different aircraft. If so, the two aircraft assigned to these legs will be severely underutilized, requiring them both to spend significant amounts of time (almost an entire day) on the ground at the spoke station. Long ground times can be eliminated, and aircraft utilization and productivity can be enhanced, by delaying the flight-leg departing the spoke by 5 to 10 minutes. This could allow sufficient time for a *single* aircraft to perform both the inbound and outbound legs. While a planner can readily recognize the value of a single retiming such as this, an optimization approach is required to track its networkwide effects and to determine the optimal set of retimings.

Rexing et al. (2000) show that allowing scheduled departures to be retimed within 20-minute intervals results in significantly improved fleetings, saving \$20 to \$50 million annually for one major airline. Interestingly, the economic gains come both from increased revenues achieved through a reduction in spill costs, and increased aircraft productivity. The retimings provide more fleet options and, hence, achieve a better overall match between flight-leg demand and assigned capacity.

Incorporating into one core problem some of the key elements of another core problem involves both *art and science*, balancing the types of information and the level of detail necessary to maintain tractability

and achieve the maximum impact on solution quality. While these enhancement efforts and the core-model mergers do not integrate the schedule design and fleet assignment problems fully, they do provide additional options and better coordination between decisions in the sequential solution process. The results are improved resource utilization and reduced costs, leading to vastly improved solutions.

2.3. Modeling for Solvability

Advanced fleet assignment, schedule design, and integrated schedule-planning models, by their nature, often take the form of network design models, with one set of decision variables assigning supply to the network and another assigning demand to the network. An example of this is the integrated *fleet assignment and passenger flow problem*, involving two types of decision variables: the assignment of aircraft types (the *supply* of seats) and the assignment of passengers to flight legs (the *demand* for seats). In basic fleet assignment models, revenue associated with demand and costs associated with supply are bundled together in the cost coefficient for each decision variable. In the integrated model, revenues are associated with demand variables and costs are associated with supply variables. This separation of costs and revenues leads to fractional aircraft decisions in the linear programming (LP) relaxation of the integrated model. This results in highly fractional solutions and large branch-and-bound trees. The weak LP relaxations of the integrated model pose major solution challenges.

One approach to resolving these challenges is to exploit the fact that there are many different, yet *equivalent*, formulations for mixed-integer problems. Equivalent formulations contain the same set of *optimal* solutions, but possibly very different sets of solutions to their associated LP relaxations. This difference is relevant to a branch-and-bound solution process: The LP bounds associated with equivalent models can vary—often dramatically.

Modeling to achieve tight relaxations is a rich research area, with substantial literature (see for example, Vanderbeck 2000 and Martin 1999). In schedule planning, research has demonstrated that models with tight LP relaxations can be achieved through an *expansion in the set of decisions represented by a variable*.

Even without the consideration of LP bounds, variable expansion has been used extensively in modeling airline-scheduling problems. In modeling core and integrated problems, expanded variables allow complex, nonlinear constraints and costs to be modeled exactly using linear constraints and a linear objective function. A case in point is the crew-pairing problem in which complex constraints are implicitly satisfied by each pairing variable, and the nonlinear crew cost structure can be represented as a linear function of the pairing variables. The fact that variable expansion can also lead to improved LP bounds—and hence improved tractability—is an added benefit.

Determining which decisions to group together within a single variable, however, is a tremendous challenge. Finding the most appropriate variable definition, capturing the complexities of the problem while allowing tractability, is very much an art. At one extreme, variables can be associated with the most elemental of decisions, such as the assignment of a crew to a flight leg, or an aircraft type to a flight leg. Such models typically have relatively small numbers of decision variables, but often suffer from factors such as weak LP relaxations, nonlinearities, and very large sets of constraints. At the other extreme, a variable can be defined that encompasses *all* decisions. The model is then structurally and algorithmically trivial—it contains one constraint requiring the greedy selection of the best solution. The difficulty with this model is, of course, its enormous number of variables. In solving scheduling problems, researchers have experimented with various expanded variable definitions to find the right balance among these alternatives.

An example can be found in the fleet assignment and passenger mix model developed by Barnhart et al. (2002). Instead of defining a fleet assignment variable as an assignment of *an aircraft type* to a *flight leg*, they expand the definition to an assignment of aircraft *types* to a *subnetwork* of flight legs. Although the expansion is motivated by the need to capture spill costs accurately, they prove that another effect of the variable expansion is that the resulting LP relaxation can provide better bounds. They illustrate, with a problem instance drawn from one major airline, that the gap between the initial LP solution and the optimal solution was just 0.3% of that for models without variable expansion. Similar findings are reported

by Armacost et al. (2002) for an application involving airline schedule design for express parcel delivery, and by Cohn and Barnhart (2003) for a service parts logistics application.

This tightening of the LP relaxation bound comes at the cost, however, of increased formulation size, with many more decision variables in the expanded-variable formulation. The effect, then, is that the expanded-variable formulation can enhance tractability with its associated improved LP relaxations, but at the same time can hamper tractability with its significantly increased model sizes. These same effects are observed in using expanded variables to model the integrated problems of aircraft routing and crew scheduling (Cohn and Barnhart 2003) and fleet assignment and schedule design (Lohatepanont and Barnhart 2001, Rexing et al. 2000).

3. Solving Scheduling Problems

As seen in §2, careful model selection can reduce complex airline planning problems to classical, well-known problems (sometimes with added side constraints) such as *set partitioning* (crew scheduling), *multicommodity flow* (fleet assignment), or *network design* (flight schedule design). These classical problems have been extensively studied and are well understood, with a sizeable body of associated literature. Even so, airline schedule-planning models present formidable solution challenges, primarily because of their immense size.

Given this, schedule-planning solution approaches prior to the 1990s focused on heuristic methods to achieve solutions. Many of the heuristics involved finding a subset of variables that could produce a quality solution, without requiring all variables to be considered explicitly. Examples of this strategy can be found in Anbil et al. (1991), Gershkoff (1989), and Hoffman and Padberg (1993). While producing improved solutions compared to earlier results, these heuristic strategies lacked the ability to indicate how much was lost in solving the models heuristically rather than exactly.

In the 1990s, with advances in computing and optimization, and motivated by the observation that even slight improvements in aircraft utilization and crew efficiency often produced significant cost savings, researchers shifted their focus to

optimization-based approaches to generate provably near-optimal solutions.

3.1. Problem-Size Reduction Methods

To manage problem size, the primary hurdle in solving airline planning models, *problem-size reduction methods* have proven essential. These techniques diminish problem size through *variable elimination*, *exploitation of dominance*, and *variable disaggregation*.

Variable Elimination. Recall that the network representation of airline flight schedules gives rise to two variable types in fleet assignment models, those corresponding to the assignment of aircraft to ground arcs and those to flight arcs. Because aircraft assignments must satisfy conservation of flow, one constraint is needed for each network node. Many of these constraints, however, are redundant. Consider the scenario, as is often the case when flights are scheduled into banks at hubs, in which an uninterrupted set of arrivals is followed by an uninterrupted set of departures. Clearly, any aircraft landing during the arrival bank will remain on the ground at least until the first departure in the departure bank. This implies that the decision variables associated with the assignment of aircraft to ground arcs within the arrival bank are unnecessary; they are fully specified by the flight-leg assignment variables corresponding to the arrival bank. Removal of the superfluous ground arc decision variables, referred to as *node consolidation*, is just one example of variable elimination techniques used in modeling and solving airline problems. Related techniques, including *arc consolidation* and *islands*, have been identified in Hane et al. (1995) and Rexing et al. (2000). Applying such reduction techniques to an extended fleet assignment model, Rexing et al. (2000) reduced their model size by more than 40%, allowing formerly intractable problem instances to be solved.

Dominance. The effectiveness of solution algorithms often depends on their ability to eliminate *dominated* variables. Shortest-path algorithms, for example, eliminate from consideration all subpaths to a node except the shortest one. Like solution algorithms, the tractability of optimization models is inextricably linked to the degree to which the principles of dominance are exploited in formulating the problem.

For example, in modeling the integrated problem of aircraft routing and crew scheduling, Cohn and Barnhart (2003) defined variables representing *complete, feasible solutions to the maintenance-routing problem*, rather than strings representing partial aircraft routings. By doing so, they replace the set of maintenance-routing constraints with a single constraint requiring the selection of a single maintenance-routing solution. In addition to eliminating a large block of constraints, their approach results in an LP relaxation that is tighter than one based on directly integrating the crew and maintenance-routing problems. The key drawback to this approach, however, is the enormous number of new variables—one per feasible maintenance-routing solution. This drawback can be overcome, though, by exploiting the fact that the only components of a maintenance-routing solution relevant to crew scheduling are the short connects it contains. (Recall that a short connect in an aircraft-routing solution is a pair of flights, assigned in succession to the same aircraft, with elapsed time between the two flights allowing sufficient time for crews to connect *only* because they stay on the same aircraft.) Short connects, when used in the crew-pairing solution, reduce crew ground time and potentially result in reduced crew costs.

The problem of finding feasible aircraft routings and minimum-cost crew solutions can therefore be thought of as finding the aircraft-routing solution and its corresponding set of short connects that allows crew costs to be minimized. Conceptualizing the problem in this way, each maintenance-routing solution can be represented by the set of *short connects* it contains. Cohn and Barnhart (2003) reported that in their experiments, a single set of short connects often represented *numerous* different aircraft-routing solutions. The result was that many of the aircraft-routing variables were redundant and could be eliminated, dramatically reducing the size of their model and allowing them to establish tractability.

The number of dominated variables is greatly influenced by variable definition. As the number of decisions encapsulated in a single variable increases, the number of variables increases and the effects of dominance become more pronounced. To illustrate this, again consider the shortest-path problem. Often, shortest-path problems are formulated using

“arc-based” variables specifying whether or not an arc is included in a shortest path. An alternate formulation based on subpaths would define variables specifying whether or not a subpath is included in the shortest path. Clearly, the subpath formulation has many more variables than the arc-based formulation. Because any subpath variable except the shortest is dominated, however, the effects of dominance are much more pronounced in the subpath formulation.

Variable Disaggregation. While *aggregated decision variables* arguably are an important mechanism for building tractable schedule-planning models, model tractability is enhanced if these aggregated variables can be *disaggregated* into variables encompassing fewer decisions.

For example, in the enhanced fleet assignment model of Barnhart et al. (2002a), decision variables correspond to subnetworks of flight legs and their assigned fleetings. The subnetworks are constructed to allow an accurate accounting of spill costs, computing spill costs as a function of itinerary flows rather than simply single flight-leg flows. The minimal subnetwork needed to compute spill costs accurately depends on the assignments of fleet types to the flight legs in that subnetwork. For example, if aircraft with relatively many seats are assigned to each flight leg, no spill results and hence spill costs can be computed exactly with subnetworks containing only a single flight leg. Thus, the minimal subnetwork size, and hence the minimal variable size needed to accurately compute spill costs, varies with the fleet assignment over the subnetwork. Barnhart et al. (2002a) use this observation in developing a variable disaggregation scheme that allowed the elimination, for one test instance, of over 90% of the variables in their enhanced fleet assignment model. This reduction in problem size allowed a previously intractable model to be solved.

3.2. Branch-and-Price Algorithms

Even with variable elimination, disaggregation, and dominance strategies, the number of remaining variables in schedule-planning models is often too large for explicit enumeration. A typical instance of the crew-pairing problem for the domestic operations of major U.S. airlines, for example, will contain *hundreds of billions* of feasible pairings. In such a case, just

the enumeration of the variables alone is intractable. Remarkably, *branch-and-price* algorithms, developed specifically for huge integer programs, can be applied successfully to solve to near-optimality even these problems in which complete variable enumeration is not practical.

Like branch-and-bound, branch-and-price is a smart enumeration strategy in which LP relaxations are solved at each node of a branch-and-bound tree to generate bounds on the optimal solution value. What makes branch-and-price particularly well suited to solving airline scheduling problems is the fact that the LPs are solved with a specialized solution approach called *column generation*. Column generation allows very large LPs to be solved without explicitly enumerating all of the variables. With advancements in both optimization techniques and computing capabilities, branch-and-price and column generation algorithms have been widely applied in the past several years, making it possible to consider very large problems in application areas such as transportation, scheduling, and combinatorial optimization (for a survey, see Barnhart et al. 1998b).

Often the strategy is to infuse the optimization process with heuristics. One commonly applied heuristic is to use column generation to solve the root node LP relaxation, and branch-and-bound to solve the *restricted* integer program (IP) containing only the subset of columns generated in solving the LP relaxation. This approach has been applied in solving the international crew-pairing problem studied by Barnhart et al. (1994) and the rostering problem considered by Ryan (1992). Although useful, and sometimes even necessary, this strategy suffers from a major shortcoming; namely, there is no guarantee that a good solution, or even a feasible solution, exists among the subset of columns producing a good LP solution.

Branch-and-price algorithms avoid this pitfall, providing a framework for finding optimal solutions to these large-scale problems. Although theoretically straightforward, the implementations of branch-and-price algorithms are nontrivial, as highlighted in the following sections.

3.2.1. Bounding: Solving the Linear Programming Relaxation. In branch-and-price, the LP relaxations are solved using *column generation*, a technique

that *implicitly*, rather than explicitly, considers all variables (often referred to as *columns*). Instead of directly solving the *master problem* (the problem with all possible columns), a *restricted master problem* containing only a subset of the original variables is solved. After solving this much smaller problem, the resulting dual information is used to identify columns with negative reduced cost to add to the restricted master problem. The restricted master problem is then re-solved, and the process repeats until no negative reduced-cost columns can be found, and the algorithm terminates with an optimal solution to the master problem.

The key steps of column generation are:

Step 1: Solve the Restricted Master Problem. Find the optimal solution to the current restricted master problem.

Step 2: Solve the Pricing Problem. Generate one or more columns with negative reduced cost. If no columns are found, STOP: The LP relaxation is solved.

Step 3: Construct a New Restricted Master Problem. Add the columns generated in Step 2 to the restricted master problem and return to Step 1.

Solution times for this algorithm are particularly sensitive to both the *number* of iterations of each step and the amount of *time* needed for each iteration. The number of iterations required in turn depends on the set of columns included in the initial restricted master problem as well as the number of columns added at each iteration. Experience shows that jump-starting the solution algorithm with a set of columns corresponding to a good heuristic solution can substantially reduce the number of required iterations. Adding many negative-reduced-cost columns at Step 3, rather than only adding the single most negative-reduced-cost column, typically leads to reductions in the number of iterations as well.

With respect to the amount of time needed to solve each step, significant gains have been made in the past few decades in reducing the times associated with Steps 1 and 3, both of which can be accomplished using standard optimization software. Lustig and Bixby (2001) detail recent advances. In contrast, solving the pricing problem (Step 2) is typically very sensitive to the models and algorithms used to solve it. Often, without tailoring the pricing-problem model and algorithm to exploit special problem structure,

even solving the root node LP with column generation can be intractable.

EXAMPLE 5: THE CREW-PAIRING PRICING PROBLEM. In the case of the crew-pairing problem, it is in fact possible to identify negative-reduced-cost variables without explicitly pricing them all out. This is done by formulating the pricing problem as a *multi-label shortest-path problem*, with minimum-cost pairings corresponding to shortest paths in an appropriately constructed network. (For a detailed exposition of multilabel shortest-path problems, see Desrochers and Soumis 1988.) Using shortest-path algorithms that exploit dominance characteristics, the shortest paths can be identified without examining all paths, and hence, the minimum reduced-cost pairings can be identified without considering all the pairings.

Although each crew pairing is represented by a network path, only a subset of paths satisfy all of the work rules and thus correspond to feasible pairings. To identify the subset of paths representing pairings, labels are used in the shortest-path algorithm. For example, a label is maintained to track the number of hours of flying in the current day's work and dynamically disallow connections to flights that violate the maximum allowed. For applications of column generation to airline crew scheduling, see Barnhart et al. (1994), Desaulniers et al. (1998), Desrosiers et al. (1991), Gamache and Soumis (1993), Lavoie et al. (1988), Ryan (1992), and Vance et al. (1997).

Even with the shortest-path framework, there are many design decisions to consider. For example, different network representations, such as arcs representing flights or arcs representing sequences of flights comprising a day of work (called a *duty period*), lead to large differences in network sizes and in the number of labels needed in the multilabel shortest-path algorithm. While more compact, networks with flight arcs (Desrosiers et al. 1995) require many more labels than duty period networks (Lavoie et al. 1988, Barnhart et al. 1994, Marsten 1994, Vance et al. 1997). In solving shortest paths in duty period networks, because only *feasible* duty periods are included, all shortest-path labels designed to ensure duty period feasibility are redundant and can be eliminated. The network representation leading to the fastest multilabel shortest-path algorithm will vary, depending on problem size and complexity of work rules.

3.2.2. Branching: Generating Integer Solutions.

The need to blend art with science extends beyond the column generation element of branch-and-price. It is also essential in developing branching strategies that maintain tractability and allow quality integer solutions to be found within reasonable run times.

A particular challenge in branch-and-price algorithms is that the standard branching rule based on variable dichotomy is often difficult to implement. The issue is that once a branching decision is taken, *all* subsequent columns generated in solving the pricing problem *must satisfy* the branching decision. This requires the pricing-problem algorithm to enforce the branching rule. If this can be done without altering the pricing algorithm, then the branching decision and the algorithm are said to be *compatible* and the efficiency of the column generation algorithm is not compromised.

Unfortunately, often the standard branching strategy based on variable dichotomy is *NOT* compatible with the pricing algorithm. Moreover, the number of branching decisions to prove optimality often is excessive for airline planning problems. To remedy this, just as in the case of column generation, researchers have embedded heuristics within the branch-and-price optimization process. An example is the *variable fixing* approach described in Marsten (1994). To reduce the number of branching decisions taken, variables with fractional values close to one are *fixed* to one sequentially. Marsten (1994) shows that by reducing the number of LPs solved in generating integer solutions, improved solutions can be generated with significantly less computing time and memory.

EXAMPLE 6: A BRANCHING RULE FOR THE CREW-PAIRING PROBLEM. Again consider the crew-pairing problem. Using branching on variable dichotomy, when a pairing variable is set to zero, it cannot be generated by the pricing algorithm, even if its reduced cost is negative. To ensure this, the multilabel shortest-path algorithm must be altered to allow for nondominated paths in the event that the dominant paths have been excluded. This alteration has deep consequences because, in the worst case, *all* pairings have to be generated, leading to an intractable pricing problem.

Instead of branching based on variable dichotomy, the branching rule most commonly used is referred to

as *branch-on-follow-ons*, derived from Ryan and Foster (1981). For crew assignment problems, the branching rule is based on pairs of flights, with one flight immediately following the other. One branching decision requires pairings containing either flight to contain *both* flights, and the other decision disallows pairings that contain both flights, with one after the other. Unlike the decisions forcing a pairing in or out of the solution, branch-on-follow-on decisions can be enforced without adversely affecting the performance of the pricing algorithm. Because the decisions are based on pairs of flights, corresponding to arcs in the network, the branching decisions can be enforced by eliminating arcs in the network. This differs from the variable dichotomy branching that affects pairings, or paths in the network. While network arcs can be easily eliminated, network paths cannot.

The branch-on-follow-on decisions are generalizable, as illustrated in Vance et al. (1994) for binary cutting-stock problems, in Savelsbergh (1997) for generalized assignment problems, in Desrochers and Soumis (1989) for urban transit crew-scheduling problems, in Desrosiers et al. (1995) and Savelsbergh and Sol (1995) for vehicle routing and scheduling problems, and in Mehrotra and Trick (1996) for graph-coloring problems.

4. Future Research and Challenges

Notwithstanding the enormous progress made in recent decades in solving the airline problems of schedule design, fleet assignment, maintenance routing, and crew scheduling, significant work remains. Research opportunities that we have identified include the following.

4.1. Core Problems

The development and application of optimization techniques to the problem of flight schedule design has immense potential, but also has equally large modeling and algorithmic challenges. More accurately, modeling the revenue function in fleet assignment problems, and more extensively adapting and applying crew-scheduling approaches to the cabin crew problem, the crew training schedule problem, and the problem of *optimizing* work rules and labor agreements, are additional, important core problem research topics that could lead to improved resource utilization.

4.2. Integrated Scheduling

Integrating crew and fleet decisions, influencing schedule decisions based on aircraft and crew resources and constraints, and expanding planning models to include other decisions and processes, such as pricing and revenue management, could all have tremendous impacts on airline resource utilization, revenue attainment, and overall profitability. In developing these integrated models, the critical challenge is to find the appropriate balance between model solvability and model scope and realism.

4.3. Robust Planning and Plan Implementation

Thus far, the majority of airline research has focused on developing efficient plans. Even a fully integrated plan, however, can be expected to be suboptimal in practice. This is because plans are based on input data that is assumed to be known and fixed. Weather conditions, mechanical problems, personnel delays and absences, and a host of other unplanned events can lead to disruptions. Moreover, because resources are in most cases highly utilized and interdependent, substantial “snowballing” effects can result from even the smallest disruptions. A delayed flight leg not only impacts subsequent legs scheduled to use the same aircraft, but future flight legs scheduled to use that flight’s crew members and passengers on affected flight legs. Thus, one disruption can lead to many downstream disruptions.

Nevertheless, airline plans are typically developed so as to maximize resource utilization. The cost of aircraft implicitly suggests the value of having minimal “ground time” between flights. Similarly, flight crews are paid such that excess time between flights increases costs. An “optimal” airline plan, therefore, is typically one with very limited slack. This lack of slack makes it difficult for disruption to be absorbed in the schedule and limits the number of options for recovery.

It is therefore reasonable to ask “Are optimal plans optimal in practice?” A number of researchers have begun to consider this question, so far with limited results. Before answering this question and identifying the *most robust* minimum-cost solution to a planning problem (that is, the solution that will perform with least cost operationally), researchers must determine the best strategy for quantifying

schedule robustness and comparing two plans. This often involves capturing how airlines recover from disruptions—another important question facing both the industry and the OR community.

Given its complexity and size, robust planning is a problem rich in opportunity and potential impact. Most successful work to date focuses on isolating causes of disruption and/or downstream effects, and incorporating within the objective function the goal of decreasing such conditions. One example is the work of Ageeva (2000), who modifies the fleet assignment problem to reward opportunities for swapping planes. Such opportunities can help to localize the impact of a disruption caused by an unavailable aircraft. Another approach to developing more robust fleet assignments is seen in Rosenberger et al. (2001a). They seek to develop solutions containing many short cycles of flights, allowing an airline to limit the number of flights cancelled when a cancellation is necessary. In crew scheduling, Klabjan and Chebalov (2003), similar to Ageeva select crew pairings with crew-swapping opportunities. Schaefer et al. (2001) and Yen and Birge (2001) take a different approach and use stochastic programming to seek more robust crew schedules.

4.4. Operations Recovery

Closely related to the issue of robustness is the airline schedule recovery problem. Given a plan and one or more disruptions, the question is how to recover in an optimal manner. Such problems are, of course, particularly challenging, involving all resources (aircraft, crew, passengers, etc.), and requiring a global view of the system. Moreover, recovery decisions often need to be made in a matter of minutes. Nonetheless, the most technologically advanced airlines have implemented the information infrastructure needed to track all resources and accommodate real-time decision making. This infrastructure, together with progress made to date, suggests that with additional research, the optimization community will be able to deliver fast, high-quality solutions consistently for the integrated aircraft, crew, and passenger recovery problem.

A number of researchers have made initial attempts at addressing this problem, typically considering only a subset of the resources involved in the disruption. Jarrah et al. (1993) consider those schedule disruptions caused by aircraft shortages. Stojkovic et al.

(2002) propose a model that allows for flight delays but not cancellations. Thengvall et al. (2000) propose a multicommodity flow model that permits both delay and cancellation decisions for aircraft recovery. Clarke (1997) proposes a similar multicommodity flow model that focuses on ensuring crew and aircraft maintenance feasible. Rosenberger et al. (2001b) propose an optimization model for aircraft recovery that partially captures passenger connections as well. Yu et al. (2003) have developed an approach to generate near-optimal crew recovery solutions.

Notwithstanding this important body of literature, much important research still remains to be done in this area, as the research community struggles to address the needs of passengers, crews, and aircraft simultaneously, while rapidly determining low-cost solutions.

4.5. Operations Paradigms

Just as deregulation in the 1970s led to significant changes in airline practices (for example, revenue management), the industry continues to face upheavals that necessitate reexamination. Even in the past 12–18 months, major events such as the terrorist attacks of September 11, 2001, and the recent economic downturn, have a wide array of impacts. Changes in security procedures, increasing costs, changing patterns of business passenger willingness to pay, and growth of low-cost competitors, all of these issues affect airline strategy and operations.

Evidence of this can be seen by recent changes to bank structures within hub-and-spoke networks. To increase aircraft and crew utilization, several airlines are considering—or have already shifted to—*depeaked* flight schedules, spreading out the arrivals and departures of aircraft at hubs over time. A depeaked schedule allows more efficient utilization of aircraft (with less time spent waiting on the ground for connecting passengers, and shorter arrival and departure queues on the ground) and spreads the workload of ground personnel and equipment more evenly throughout the day, thereby eliminating the need for additional resources during peak periods. The disadvantage of this strategy is less coordinated, longer average passenger connection times at hubs. Related research questions involve restructuring of the revenue model to account for schedule performance, both planned

and operational. Optimization is playing, and can continue to play, an important role in assessing the systemwide cost and service impacts of this, and numerous other, operating paradigm shifts.

References

- Abara, J. 1989. Applying integer linear programming to the fleet assignment problem. *Interfaces* **19**(4) 20–28.
- Ageeva, Y. 2000. Approaches to incorporating robustness into airline scheduling. MIT thesis, Cambridge, MA.
- Anbil, R., E. Gelman, B. Patty, R. Tanga. 1991. Recent advances in crew-pairing optimization at American Airlines. *Interfaces* **21** 62–74.
- Armacost, A., C. Barnhart, K. Ware. 2002. Composite variable formulations for express shipment service network design. *Transportation Sci.* **36**(1) 1–20.
- Barnhart, C., A. Farahat, M. Lohatepanont. 2002a. Airline fleet assignment with enhanced revenue modeling: An alternative model and solution approach. Working paper, Operations Research Center, MIT, Cambridge, MA.
- Barnhart, C., T. Kniker, M. Lohatepanont. 2002b. Itinerary-based airline fleet assignment. *Transportation Sci.* **36**(2) 199–217.
- Barnhart, C., F. Lu, R. Sheno. 1998a. Integrated airline scheduling. G. Yu, ed. *Operations Research in the Air Industry*. International Series in *Operations Research and Management Science*, Vol. 9, 384–403.
- Barnhart, C., E. Johnson, R. Anbil, L. Hatay. 1994. A column generation technique for the long-haul crew assignment problem. T. Ciriano, R. Leachman, eds. *Optimization in Industry: Vol. II*. Wiley, New York, 7–22.
- Barnhart, C., E. Johnson, G. Nemhauser, M. Savelsbergh, P. Vance. 1998b. Branch-and-price: Column generation for solving huge integer programs. *Oper. Res.* **46**(3) 316–329.
- Barnhart, C., N. Boland, L. Clarke, E. Johnson, G. Nemhauser, R. Sheno. 1998c. Flight string models for aircraft fleet and routing. *Transportation Sci.* **32**(3) 208–220.
- Barnhart, C., A. M. Cohn, E. L. Johnson, D. Klabjan, G. L. Nemhauser, P. H. Vance. 2003. Airline crew scheduling. Randolph W. Hall, ed. *Handbook of Transportation Science*, 2nd ed. Kluwer Academic Publishers, Norwell, MA.
- Butchers, E. R., P. R. Day, A. P. Goldie, S. Miller, J. A. Meyer, D. M. Ryan, A. C. Scott, C. A. Wallace. 2001. Optimized crew scheduling at Air New Zealand. *Interfaces* **31** 30–56.
- Cappanera, P., G. Gallo. 2001. On the airline crew rostering problem. Technical Report TR-01-08, Department of Computer Science, University of Pisa, Italy.
- Caprara, A., P. Toth, M. Fischetti, D. Vigo. 1998. Modeling and solving the crew rostering problem. *Oper. Res.* **46** 820–830.
- Christou, I. T., A. Zakarian, J. Liu, H. Carter. 1999. A two-phase genetic algorithm for large-scale bidline-generation problems at Delta Air Lines. *Interfaces* **29** 51–65.
- Clarke, L., C. Hane, E. Johnson, G. Nemhauser. 1996a. Maintenance and crew considerations in fleet assignment. *Transportation Sci.* **30**(3) 249–260.
- Clarke, L., E. Johnson, G. Nemhauser, Z. Zhu. 1996b. The aircraft rotation problem. *Ann. Oper. Res.* **69** 33–46.

- Clarke, M. 1997. Development of heuristic procedures for flight rescheduling in the aftermath of irregular operations. Working paper, Department of Aeronautical/Astronautical Engineering, MIT, Cambridge, MA.
- Clarke, M., B. Smith. 2000. The impact of operations research on the evolution of the airline industry: A review of the airline planning process. Research paper, Sabre Inc., Dallas, TX.
- Cohn, A., C. Barnhart. 2002. Composite-variable modeling for service parts logistics. Submitted to *Ann. Oper. Res.*
- Cohn, A., C. Barnhart. 2003. Improving crew scheduling by incorporating key maintenance routing decisions. *Oper. Res.* **51** 387–396.
- Cook, T. 2000. Creating competitive advantage using model-driven support systems. Presentation in the *MIT Global Airline Indust. Study Distinguished Speaker Sem. Ser.*, Cambridge, MA.
- Cordeau, J., G. Stojkovic, F. Soumis, J. Desrosiers. 2000. Benders decomposition for simultaneous aircraft routing and crew scheduling. Technical Report G-2000-37, GERAD, École Polytechnique de Montréal, Québec, Canada.
- Dawid, H., J. König, C. Strauss. 2001. An enhanced rostering model for airline crews. *Comput. Oper. Res.* **28** 671–688.
- Day, P. R., D. M. Ryan. 1997. Flight attendant rostering for short-haul airline operations. *Oper. Res.* **45** 649–661.
- Desaulniers, G., J. Desrosiers, M. Gamache, F. Soumis. 1998. Crew scheduling in air transportation. T. Crainic, G. Laporte, eds. *Fleet Management and Logistics*. Kluwer Academic Publishers, Boston, MA, 169–185.
- Desrochers, M., F. Soumis. 1988. A generalized permanent labeling algorithm for the shortest path problem with time windows. *INFOR* **26** 191–212.
- Desrochers, M., F. Soumis. 1989. A column generation approach to the urban transit crew scheduling problem. *Transportation Sci.* **23** 1–13.
- Desrosiers, J., Y. Dumas, M. M. Solomon, F. Soumis. 1995. Time constrained routing and scheduling. *Vehicle Routing*, North-Holland, Amsterdam.
- Desrosiers, J., Y. Dumas, M. Desrochers, F. Soumis, B. Sanso, P. Trudeau. 1991. A breakthrough in airline crew scheduling. Report G-91-11, GERAD, Montréal, Canada.
- Erdmann, A., A. Noltemeier, R. Schrader. 1999. Modeling and solving the airline schedule generation problem. Technical Report zpr99-351, ZAIK, University of Cologne, Germany.
- Gamache, M., F. Soumis. 1993. A method for optimally solving the rostering problem. Les Cahier du GERAD, G-90-40, École des Hautes Études Commerciales, Montréal, Canada.
- Gamache, M., F. Soumis, D. Villeneuve, J. Desrosiers, E. Gélinas. 1998. The preferential bidding system at Air Canada. *Transportation Sci.* **32** 246–255.
- Gershkoff, I. 1989. Optimizing flight crew schedules. *Interfaces* **19** 29–43.
- Gopalan, R., K. Talluri. 1998. The aircraft maintenance routing problem. *Oper. Res.* **46** 260–271.
- Hane, C. A., C. Barnhart, E. L. Johnson, R. E. Marsten, G. L. Nemhauser, G. Sigismondi. 1995. The fleet assignment problem: Solving a large-scale integer program. *Math. Programming* **70** 211–232.
- Hoffman, K. L., M. Padberg. 1993. Solving airline crew-scheduling problems by branch-and-cut. *Management Sci.* **39** 657–682.
- Jacobs, T. L., E. L. Johnson, B. C. Smith. 1999. O&D FAM: Incorporating passenger flows into the fleet process. *AGIFORS Sympos.*, New Orleans, LA.
- Jarrah, A., G. Yu, N. Krishnamurthy, A. Rakshit. 1993. A decision support framework for airline flight cancellations and delays. *Transportation Sci.* **27**(3) 266–280.
- Klabjan, J., S. Chebalov. 2003. Robust airline crew pairing optimization. *FAA-NEXTOR-INFORMS Sympos.*, Falls Church, VA.
- Kohl, N., S. E. Karisch. 2002. Airline crew rostering: Problem types, modeling, and optimization. Working paper, Carmen Consulting, Copenhagen, Denmark.
- Lavoie, S., M. Minoux, E. Odier. 1988. A new approach for crew pairing problems by column generation with an application to air transportation. *Eur. J. Oper. Res.* **35** 45–58.
- Lohatepanont, M., C. Barnhart. 2001. Airline schedule planning: Integrated models and algorithms for schedule design and fleet assignment. *Transportation Sci.* Forthcoming.
- Lustig, I., R. Bixby. 2001. A historical perspective on linear programming performance. http://www.ilog.com/products/optimization/times_winter2001/tech_historical.cfm.
- Marsten, R. 1994. Crew planning at Delta Airlines. *Math. Programming Sympos. XV*, Ann Arbor, MI.
- Martin, R. C. 1999. *Large Scale Linear and Integer Optimization: A Unified Approach*. Kluwer Academic Press, Norwell, MA.
- Mehrotra, A., M. A. Trick. 1996. A column generation approach for exact graph coloring. *INFORMS J. Comput.* **8** 344–354.
- Patty, B., E. Gelman, R. Tanga, R. Anbil. 1991. Recent advances in crew-pairing optimization at American Airlines. *Interfaces* **21** 62–74.
- Prousaloglou, K. E., F. S. Koppelman. 1995. Air carrier demand: An analysis of market share determinants. *Transportation* **22** 371–388.
- Prousaloglou, K. E., F. S. Koppelman. 1999. The choice of carrier, flight and fare class. *J. Air Transport Management* **5**(4) 193–201.
- Rexing, B., C. Barnhart, T. Kniker, A. Jarrah, N. Krishnamurthy. 2000. Airline fleet assignment with time windows. *Transportation Sci.* **34**(1) 1–20.
- Rosenberger, J., E. Johnson, G. Nemhauser. 2001a. A robust fleet assignment model with hub isolation and short cycles. Working paper, SISE, Georgia Institute of Technology, Atlanta, GA.
- Rosenberger, J., E. Johnson, G. Nemhauser. 2001b. Rerouting aircraft for airline recovery. Working paper, SISE, Georgia Institute of Technology, Atlanta, GA.
- Rushmeier, R., S. Kontogiorgis. 1997. Advances in the optimization of airline fleet assignment. *Transportation Sci.* **31**(2) 159–169.
- Ryan, D. M. 1992. The solution of massive generalized set partitioning problems in air crew rostering. *J. Oper. Res. Soc.* **43** 459–467.
- Ryan, D. M., B. A. Foster. 1981. An integer programming approach to scheduling. A. Wren, ed. *Computer Scheduling of Public Transport Urban Passenger Vehicle and Crew Scheduling*. North-Holland, Amsterdam, The Netherlands, 269–280.
- Savelsbergh, M. W. P. 1997. A branch-and-price algorithm for the generalized assignment problem. *Oper. Res.* **45** 831–841.
- Savelsbergh, M. W. P., M. Sol. 1995. The general pickup and delivery problem. *Transportation Sci.* **29** 17–29.
- Schaefer, A., E. Johnson, A. Kleywegt, G. Nemhauser. 2001. Airline crew scheduling under uncertainty. Working paper, SISE, Georgia Institute of Technology, Atlanta, GA.
- Smith, B. C., J. F. Leimkuhler, R. M. Darrow. 1992. Yield management at American Airlines. *Interfaces* **22** 8–31.
- Stojkovic, G., F. Soumis, J. Desrosiers, M. Solomon. 2002. An optimization model for a real-time flight scheduling problem. *Transportation Res.* **36A** 779–788.

- Thengvall, B., G. Yu, J. Bard. 2000. Balancing user preferences for aircraft schedule recovery during irregular airline operations. *IIE Trans.* **32** 181–193.
- Vance, P. H., C. Barnhart, E. L. Johnson, G. L. Nemhauser. 1994. Solving binary cutting stock problems by column generation and branch-and-bound. *Comput. Optim. Appl.* **3** 111–130.
- Vance, P. H., C. Barnhart, E. L. Johnson, G. L. Nemhauser. 1997. Airline crew scheduling: A new formulation and decomposition algorithm. *Oper. Res.* **45** 188–200.
- Vanderbeck, F. 2000. On Dantzig-Wolfe decomposition in integer programming and ways to perform branching in a branch-and-price algorithm. *Oper. Res.* **48** 111–128.
- Wiper, D. S., J. D. Quillinan, R. Subramanian, R. P. Scheff, Jr., R. E. Marsten. 1994. Coldstart: Fleet assignment at Delta Air Lines. *Interfaces* **24** 104–120.
- Yen, J., J. Birge. 2001. A stochastic programming approach to the airline crew scheduling problem. Submitted to *Oper. Res.*
- Yu, G., M. Argüello, G. Song, S. McCowan, A. White. 2003. A new era for crew recovery at Continental Airlines. *Interfaces* **33** 5–22.