Airline Crew Scheduling under Uncertainty

Andrew J. Schaefer *

Department of Industrial Engineering

University of Pittsburgh

Pittsburgh, Pa. 15261

Ellis L. Johnson

Anton J. Kleywegt

George L. Nemhauser

School of Industrial and Systems Engineering

Georgia Institute of Technology

Atlanta, Ga. 30332

May 28, 2001

Abstract

Airline crew scheduling algorithms widely used in practice assume no disruptions. Since disruptions often occur, the actual cost of the resulting crew schedules is often significantly greater. We consider algorithms for finding crew schedules that perform well in practice. The deterministic crew scheduling model is an approximation of crew scheduling under uncertainty under the assumption that all pairings will operate as planned. We seek better approximate solution methods for crew scheduling under un-

*Corresponding author. schaefer@engrng.pitt.edu

certainty that still remain tractable. We give computational results from three fleets that indicate that the crew schedules obtained from our methodology perform better in operations than the crew schedules found via state-of-the-art methods. We provide a lower bound on the cost of an optimal crew schedule in operations, and demonstrate that some of the crew schedules found using our methodology perform very well relative to this lower bound.

For major domestic carriers crew costs are second only to fuel costs, and can exceed a billion dollars annually. Therefore, airlines devote great effort to planning good crew schedules. But the planning problem can be very difficult to solve because there are many governmental and contractual regulations concerning pilots, and problems found in practice often have billions of possible solutions.

There is currently a great deal of concern about air traffic congestion. In June 2000, flight delays were up over 16% from June 1999 Phillips and Irwin, 2000. Moreover, air traffic in America and Europe is expected to double in the next 10-15 years. If airport capacity remains constant, it is estimated that each 1% increase in airport traffic will bring about a 5% increase in delays Anonymous, 2000.

Disruptions are becoming more frequent and more severe. The Air Transport Association reports that the average daily number of delays of more than 15 minutes tied to air-traffic control increased to 2,149 in 1999, up from 1,807 in 1998 and 1,416 in 1997. They also estimate that delays cost consumers and airlines about $5.2 billion in 1999, up from $5 billion the year before Mathews, 2000. The Federal Aviation Administration (FAA) reported a 58% increase in delays from 1995 to 1999, and flight cancellations increased 68% over the same period. Atlanta's Hartsfield International Airport had a 138% increase in flight cancellations. The total cost to the airlines alone at Hartsfield was estimated at $250.9 million in 1999 O'Dell, 2000.

Crew planning affects airline operations. Airlines must delay flights if a crew is unavailable. Recently, pilots for some major domestic carriers refused to work overtime to protest the pace of negotiations of a new contract. The resulting shortage of pilots forced numerous cancellations.

Although airlines operate in a highly uncertain environment, very few airline planning models consider uncertainty in operations. Some of this is due to the structure of airline management. A plan is typically evaluated not by operational performance but by the quality of the plan *assuming it may be implemented in* *operations*. With the exception of yield management, we know of no airline planning models that measure the quality of a plan by its performance in operations. The integration of airline planning and operations is a fertile area of great practical and theoretical interest.

We define two classes of airline disruptions based on the length of the disruption. A *frictional* disruption is of limited duration. Examples include delays due to connecting passengers, airport congestion, brief unscheduled maintenance incidents, and localized, mild weather systems. The other class of disruptions is *severe*, which include lengthy unscheduled maintenance disruptions, and large-scale severe weather systems. This classification is not a strict dichotomy; disruptions may have aspects of both frictional and severe disruptions. We limit this study to planning and operations under frictional disruptions. We discuss the state-of-the-art in airline operations in Section 1.2. We show that under frictional delays, the state-of-the-art crew scheduling model can be improved not only in terms of pilot compensation but in terms of other measures such as on-time performance.

1.1 Crew Scheduling

Typically pilots may only fly one type of aircraft. Therefore the crew scheduling problem is separable by fleet type. When a crew is on duty, it flies a set of consecutive flight legs that follow certain regulations and contractual restrictions. Such a set of legs is called a *duty*. The *sit time* is the time between two consecutive legs within a duty. The number of minutes that elapse between the beginning of a duty and the end of the duty is the *elapsed time*. The elapsed time includes a briefing period before the first leg of the duty, and a debriefing period after the last leg of the duty.

A *pairing* or crew trip is a set of duties. Consecutive duties must be separated by a *rest* period. A pairing must begin and end at a specified station; such stations are called *crew bases*. Pairings flown within the U.S. must adhere to certain FAA as well as contractual rules. For instance, one rule requires that a crew that flies more than 8 hours within a 24 hour period must receive compensatory rest FAA, 1999. The *time away from base* (TAFB) of a pairing is the number of minutes that elapse between the beginning of the pairing and the end of the pairing. In many instances, crews are paid based on the amount of time they fly in their pairing. However, there is a minimum guaranteed pay for any pairing, and there is additional compensation for the crew if the TAFB of the pairing or the elapsed time of one or more of the duties is large enough. We describe the details of calculating crew cost in Section 2.

Since a crew can fly only one fleet type, the fleet assignment problem and the aircraft rotation problem are typically solved before the crew scheduling problem. If a crew flies two consecutive legs on different planes, the scheduled connection time between these legs must exceed a *minimum connection time*. However, if the crew remains on the same plane for two consecutive flights, there is no minimum connection time.

A *crew schedule* is a set of pairings that partitions the legs to be flown by a single fleet. Crew scheduling problems are solved by generating pairings and solving an integer program. The *daily crew scheduling problem* is solved under the assumption that each leg is flown every day.

The crew scheduling problem is usually modeled as a set partitioning problem

$$\{\min\ cx\ :\ Ax\ =\ 1, x \text{ binary}\} \tag{1}$$

where a_{ij}, the ij^{th} entry of the matrix A, is 1 if pairing j flies leg i, and 0 otherwise.

There may be a large number of pairings for a relatively small fleet. Vance et al. 1997 found that a fleet with 250 daily legs had over 5,000,000 pairings. Larger fleets have billions of legal pairings. The enormous number of pairings is a major difficulty in solving airline crew scheduling problems exactly.

Recent work on deterministic airline crew scheduling include Lavoie et al. 1988, Gershkoff 1989, Anbil et al. (1991, 1992a, 1992b, 1999), Graves et al. 1993, Hoffman and Padberg 1993, Barnhart et al. 1994, Andersson et al. 1997, Chu et al. 1997, Vance et al. 1997, Klabjan et al. (1999a, 1999b, 1999c) and Snowden et al. 2000. Yen 2000 also considers the problem of crew scheduling under uncertainty. She formulates the problem as a two-stage stochastic program, where the first stage is the crew scheduling problem and the second stage involves penalties for delays. She does not consider the operational cost of a crew schedule, and provides computational results for a few fleets smaller than fleets flown by major carriers.

1.2 Airline Operations

Recovery is the process of reacting to a disruption. The optimal recovery decision is hard to determine. The future is uncertain, and canceling a leg or rerouting a crew or a plane can have costly consequences throughout the airline's system. In practice, airlines make recovery decisions manually with little decision support Lettovský, 1997. This makes airline recovery difficult to model because airline decision-makers use intuition and subjective judgment. Most optimization research done on airline operations has been on crew recovery, but these models assume all legs will be flown according to their new scheduled leg times. We are unaware of any research on dynamic and stochastic airline recovery models. In many respects, finding an optimal recovery policy is more challenging than the airline planning process. It is usually acceptable to solve planning problems using algorithms that may take many hours to run since the plans are made many months ahead of time. While airline planning models may be solved sequentially with a long time between decision phases, problems such as crew rescheduling, plane and passenger rerouting must be solved rapidly.

Recovery is very important to an airline because costs may be high during disrupted operations. Lettovský 1997 reports that irregular operations can be responsible for as much as 3% of an airline's operating expenses. Caldwell 1997 reported that in 1992 Delta Airlines had irregular operations that disrupted 8.5 million passengers, and cost Delta up to $500 million in direct costs, excluding any loss of goodwill. There has been very little research done towards making optimal recovery decisions. In practice, recovery decisions are made in the *Airline Operations Control Center* (AOC). Various coordinators work together to find real-time recovery solutions when an airline is faced with a disruption. An AOC must consider planes, pilots, passengers, flight attendants and cargo. Many options are available to an AOC in recovery. Flights may be cancelled or delayed. Planes may be diverted or *ferried* to a destination without passengers. Crews may be rescheduled. Reserve crews may be called. Flight attendants may be rerouted, and reserve flight attendants may be called. Passengers may be rescheduled, and may fly on other airlines.

Integrated recovery models simultaneously consider crew, aircraft rotation, and passenger recovery problems. Integrated recovery models also determine how long to delay a leg or whether to cancel a leg. They reroute planes, passengers, and crews. Lettovský 1997 describes an integrated recovery model that uses Benders' decomposition 1962 for crew, aircraft, and passenger recovery. The integrated recovery master problem is a fleet assignment model with a limit on the number of arrivals in a given time period at a station. No implementation details or computational results are provided with the Lettovský model.

In Section 2 we discuss methods of evaluating the quality of a crew schedule. In Section 3 we describe SimAir, a stochastic simulation of airline operations. In Section 4 we give two algorithms for finding crew schedules that may perform well in operations. We also provide a method of finding a lower bound on the expected operational cost of a crew schedule. In Section 5 we provide computational results for three fleets from a major domestic carrier.

2 Evaluating a Crew Schedule

While finding good crew schedules is critical for airlines, an important question is what is meant by a "good" schedule. Airlines have traditionally evaluated a crew schedule by its planned cost. This implicitly assumes that every leg will be operated as planned, but anecdotal evidence suggests that this hardly ever happens. We propose the evaluation of a crew schedule by its performance in operations. To evaluate a crew schedule's performance in operations, we must first specify mechanisms and probabilities of disruptions, as well as a recovery policy. In order to find a best crew schedule, we must prescribe a method of comparing two different crew schedules.

Crew schedules can affect pilot compensation and on-time performance. The Bureau of Transportation Statistics (BTS) defines a leg to be on-time if it arrives no later than 15 minutes after its scheduled arrival time BTS, 1998. A poor performance in these rankings can adversely affect the public's perception of an airline.

Pilot compensation is the second-largest direct cost incurred by major domestic airlines. Only fuel is more costly. Hoffman and Padberg 1993 reported that total pilot compensation exceeded $1.4 billion annually at the largest domestic airlines, and a senior pilot earned up to $250,000 annually. These figures are larger now.

There are two ways of measuring pilot compensation. The planned cost of a crew schedule is the sum of the planned pairing costs over all pairings in the schedule, where planned pairing costs are given by a closed-form expression. The planned cost of a crew schedule is widely used by domestic airlines to evaluate crew schedules. The operational cost of a crew schedule is the pilots' compensation under operations. Since operational conditions are not known at the planning stage, the operational cost of a crew schedule is a random variable. Therefore, it is not clear how to compare the operational costs of two different crew schedules.

We discuss the operational cost of a crew schedule in Section 2.2. We address methods for comparing different crew schedules in Section 3. We denote planned, deterministic quantities by underlining them, and operational, unknown quantities by placing a tilde over them. For example, the planned flight time, or *block time* of leg l is denoted by $\underline{block}(l)$ and the operational block time of leg l is denoted by $\widetilde{block}(l)$. When a quantity is the same in planning and operations, it appears with neither. Thus, for instance, the length of the briefing period is given by $brief$, since it is modeled to be the same in operations and planning.

2.1 The Planned Cost of a Crew Schedule

The airline industry does not measure the cost of a crew schedule in monetary terms. Rather, it is expressed in terms of minutes of pay and credit. When crew schedules are found, pairings are not yet assigned to particular pilots. Since pilot salaries differ, determining the monetary cost of a crew schedule is only possible once pairings have been assigned to particular pilots. The flight-time-credit (FTC) of a duty is the difference between its total cost in minutes of pay and credit and the total block time expressed as a percentage of the total block time of the duty. A similar measure exists for pairings and crew schedules. We will let $\underline{FTC}(\cdot)$ denote the planned FTC of any duty, pairing or crew schedule. The method for calculating the planned cost of a crew schedule varies by airline. We give an example of one method.

Let q be any pairing consisting of duties $d_1, \ldots, d_k$. For $1 \leq i \leq k$, duty d_i consists of legs $l_{i,1}, \ldots, l_{i,m(i)}$. For $1 \leq j \leq m(i)$, let $\underline{dep}(l_{i,j})$ be the scheduled departure of leg $l_{i,j}$ in minutes and let $\underline{arr}(l_{i,j})$ be its scheduled arrival time in minutes. These times are relative to the start of the pairing, so that for $1 \leq i \leq k-1$,

$\underline{dep}(l_{i+1,1}) \; > \; \underline{arr}(l_{i,m(i)})$. Our convention is that $\underline{dep}(l_{1,1}) = 0$, and 1440 occurs exactly one day after the pairing has begun. Let $\underline{block}(l_{i,j})$ be the planned block time of leg $l_{i,j}$ in minutes, defined by

$$\underline{block}(l_{i,j}) \;=\; \underline{arr}(l_{i,j}) \;-\; \underline{dep}(l_{i,j}). \tag{2}$$

Let *brief* be the length of the pilot briefing period prior to every duty. Let *debrief* be the length of the pilot debriefing period after every duty. The parameters *brief* and *debrief* are constants and are in minutes. For each duty d_i, $1 \leq i \leq k$, define the planned *elapsed* time of the duty as

$$\underline{elapse}(d_i) \;=\; \underline{arr}(l_{i,m(i)}) \;-\; \underline{dep}(l_{i,1}) \;+\; brief \;+\; debrief. \tag{3}$$

Let $r_e < 1$ be a fraction representing the rate of pay for elapsed time in terms of minutes of pay and credit. Let mg_d be the minimum guarantee for a duty, which is given in minutes of pay and credit. The *planned duty cost* of duty d_i is expressed in minutes of pay and credit and is given by

$$\underline{b}(d_i) \;=\; \max\left\{\sum_{j=1}^{m(i)} \underline{block}(l_{i,j}), r_e \times \underline{elapse}(d_i), mg_d\right\}. \tag{4}$$

The planned flight-time-credit (FTC) of duty d_i is given by

$$\underline{FTC}(d_i) \;=\; \frac{\underline{b}(d_i) - \sum_{j=1}^{m(i)} \underline{block}(l_{i,j})}{\sum_{j=1}^{m(i)} \underline{block}(l_{i,j})}. \tag{5}$$

The planned *time away from base* of pairing q is the total number of minutes that elapse during the pairing given by

$$\underline{TAFB}(q) \;=\; \underline{arr}(l_{k,m(k)}) \;-\; \underline{dep}(l_{1,1}) \;+\; brief \;+\; debrief. \tag{6}$$

Let $r_t < 1$ be a fraction representing the rate of pay of time away from base. Let mg_p be a minimum guarantee per duty in a pairing. Then the *planned pairing cost* of pairing q is given by

$$\underline{c}_q = \max\left\{\sum_{i=1}^{k} \underline{b}(d_i), r_t \times \underline{TAFB}(q), mg_p \times k\right\}. \tag{7}$$

Vance et al. 1997 use values of $r_e = \frac{4}{7}$, $mg_d = 0$, $r_t = \frac{2}{7}$, and $mg_p = 300$.

The planned FTC of pairing q is defined by

$$\underline{FTC}(q) = \frac{\underline{c}_q - \sum_{i=1}^{k}\sum_{j=1}^{m(i)} \underline{block}(l_{i,j})}{\sum_{i=1}^{k}\sum_{j=1}^{m(i)} \underline{block}(l_{i,j})}. \tag{8}$$

Let $\underline{c}(C)$ be the planned cost of a crew schedule C consisting of pairings $q_1, \ldots, q_{|C|}$, given by

$$\underline{c}(C) = \sum_{q \in C} \underline{c}_q. \tag{9}$$

Let $\underline{block}(C)$ be the total scheduled block time of all legs in the flight schedule. The planned FTC of crew schedule C is

$$\underline{FTC}(C) = \frac{\underline{c}(C) - \underline{block}(C)}{\underline{block}(C)}. \tag{10}$$

2.2 The Operational Cost of a Crew Schedule

The operational cost of a schedule is the sum of the operational costs of the pairings that comprise it. Different airlines may have different methods of calculating the operational cost of a pairing. We give an example of how one major domestic carrier calculates the operational cost of a pairing.

For any leg l, let $\widetilde{dep}(l)$ be the actual departure time of the leg, and let $\widetilde{arr}(l)$ be its actual arrival time. The operational block time of the leg is defined in the same way as its planned block time, but the operational departure and arrival times are used. For any duty d the operational elapsed time is calculated in the same way as its planned elapsed time, except the actual arrival time of its last leg is used. Its operational cost is calculated in a similar way as its planned cost, but it considers the operational block times of legs and the operational elapsed time. For any pairing q the operational time away from base is the same as its planned time away from base, except the actual arrival time of its last leg is used. The operational cost of pairing q is given by

$$\widetilde{c}_q \;=\; \max\left\{\sum_{i=1}^{k}\widetilde{b}(d_i),\; r_t \times \widetilde{TAFB}(q),\; k\cdot mg_p, \underline{c}_q\right\}. \tag{11}$$

For any crew schedule C define its *operational cost* by $\widetilde{c}(C) = \sum_{q\in C}\widetilde{c}_q$. Notice that $\tilde{c}(C) \;\geq\; \underline{c}(C)$, since by (11) $\widetilde{c}_q \;\geq\; \underline{c}_q$ for all pairings $q \in C$. The operational FTCs of duties, pairings and crew schedules are calculated in the same manner as planned FTCs except the operational costs replace the planned costs. In this paper we assume that no flights are cancelled and that each duty and pairing fly the legs originally assigned to it, although possibly with different departure and arrival times.

2.3 On-Time Performance

A leg is *on-time* if it arrives no later than 15 minutes after its scheduled arrival time. The on-time performance of an airline depends on many factors. The block time is an important factor. Clearly, on-time performance improves with longer scheduled block times. The flight schedule itself may influence an airline's on-time performance. An airline that operates many legs at times and airports with much congestion may have a worse on-time performance than an airline that does not. The aircraft rotation may also affect the on-time performance. A routing in which the planes have very little time between legs may result in a lower on-time performance than a routing in which the planes have extra time between legs. In addition to all these factors, a crew schedule may affect an airline's on-time performance.

3 SimAir - A Simulation of Airline Operations

In order to find crew schedules that perform well in operations we need a method for evaluating the operational performance of a crew schedule. It is impractical to test multiple crew schedules by running them in actual operations. It may require many days to estimate accurately the expected performance of a crew schedule, and we may need to require many crew schedules. It is unlikely that any airline would allow experiments with alternate crew schedules without any indication that improved performance is likely.

Evaluating the operational performance of a crew schedule must be inexpensive and accurate.

We use SimAir, a Monte Carlo simulation of airline operations, to evaluate a crew schedule's performance. We present an abbreviated description of SimAir, which is based on Rosenberger et al. 2000a. A more detailed description of the stochastic model underlying SimAir is given in Rosenberger et al. 2000b. SimAir is a flexible simulation that permits the study of a crew schedule under a recovery method and delay distribution. SimAir explicitly considers crews, planes, and passengers. SimAir provides a method for evaluating the performance of a crew schedule in operations.

For any leg l, let ω_l be a random block time error. SimAir updates the arrival time of leg l to be

$$\widetilde{arr}(l) \;=\; \underline{block}(l) \;+\; \widetilde{dep}(l) \;+\; \omega_l, \tag{12}$$

where $\underline{block}(l)$ is the scheduled block time of leg l. Let $\widetilde{ctime}(l)$ be the earliest time when the crew is available to fly leg l. Let $\widetilde{ptime}(l)$ be the earliest time when the plane is available to fly leg l. SimAir schedules leg l to depart at time $\widetilde{sdtime}(l)$ which is defined by

$$\widetilde{sdtime}(l) \;=\; \max\{\widetilde{ctime}(l), \widetilde{ptime}(l), \underline{dep}(l)\}. \tag{13}$$

Ground time is the time from the plane and crew are available until the leg departs. Ground time delays may occur due to connecting passengers and cargo, airport congestion, and so on. Let ξ_l be a nonnegative random variable denoting the length of a ground time delay. Leg l will depart at time

$$\widetilde{dep}(l) \;=\; \widetilde{sdtime}(l) \;+\xi_l. \tag{14}$$

Although SimAir can generate delays from a variety of sources, to evaluate a crew schedule it is not necessary to explicitly consider the source of delays. Instead, we use aggregate distributions for additional block time and ground time. A block time disruption affects the number of minutes a crew flies, but a ground time disruption does not. Unscheduled maintenance is considered separately since it affects a specific plane.

When a flight is delayed, SimAir must find a recovery action that responds to the delay. SimAir may use an external recovery program, or it may use a simple routine that waits for the scheduled planes and crews regardless of their tardiness. We refer to this recovery as *push-back*. Consider the arrival of leg l_i, where the crew's next leg is l_j and the plane's next leg is l_k, where l_j is not necessarily the same as l_k. Upon the arrival of leg l_i SimAir calculates $\widetilde{ptime}(l_k)$, the plane's ready time for the plane's next flight and $\widetilde{ctime}(l_j)$, the crew's ready time for the crew's next flight. Let *minplaneturn* be the minimum amount of time a plane needs to turn. Let $\widetilde{maint}(l)$ be a random variable indicating the length of an unscheduled maintenance delay prior to the departure of leg l. The ready time of the plane upon the arrival, $\widetilde{ptime}(l_k)$ is given by

$$\widetilde{ptime}(l_k) \;=\; \widetilde{arr}(l_i) \;+\; minplaneturn \;+\; \widetilde{maint}(l_k). \tag{15}$$

The crew is ready at $\widetilde{ctime}$, which if the crew has not completed a duty is given by

$$\widetilde{ctime}(l_j) \;=\; \widetilde{arr}(l_i) \;+\; \widetilde{\theta}_{l_i}. \tag{16}$$

The random variable $\widetilde{\theta}_{l_i}$ is the amount of time needed to keep the crew connection legal between legs l_i and l_j. If l_i is not the last leg in a duty,

$$\widetilde{\theta}_{l_i} \;=\; crewminturn \;+\; (swapplane \times swaptime), \tag{17}$$

where *crewminturn* is the crew's minimum turn time if it stays with the plane, *swaptime* is the additional time needed if the crew changes planes, and

$$swapplane = \begin{cases} 1 & \text{if the crew changes planes} \\ 0 & \text{otherwise.} \end{cases}$$

Otherwise, if the crew has just completed a duty but has not yet completed its pairing, then

$$\widetilde{\theta}_{l_i} \;=\; \widetilde{minrest} \;+\; brief \;+\; debrief, \tag{18}$$

where $\widetilde{minrest}$ is the minimum amount of rest required by the recovery policy or by the relevant regulations. It may depend on the history of the pairing, and will be longer if compensatory rest is required. The random variable $\widetilde{minrest}$ depends on the legs previously flown by the crew. See Rosenberger et al. 2000b for more information.

4 Methodology

Exact formulations for crew scheduling under uncertainty are intractable due to the enormous state space, action space, and number of time periods required. The deterministic crew scheduling model is an approximation of crew scheduling under uncertainty under the assumption that all pairings will operate as planned. We seek better approximate solution methods for crew scheduling under uncertainty that still remain tractable.

Airline crew scheduling under uncertainty could be formulated as a Markov decision process (MDP). However, such a model would be intractable. The state of the system describes every aspect of the system that is relevant for operational decisions. The state must contain information about the current status and history of every crew member and plane, as well as a description of the current operating environment including weather, congestion, and so on. The first stage consists of the planning period, where a flight schedule, fleet assignment, routing, and initial crew schedule are found. Operational decisions are made in subsequent stages. The number of stages could be quite large, since the state of the system can change within minutes. The action space consists of all possible feasible decisions. These include cancelling flights, rerouting planes and passengers, rescheduling crews, and so on. These operational decisions may have a profound impact on the legality of future crews due to complicated regulations such as the 8-in-24 rule.

We introduce two methods for finding crew schedules that may perform well in operations. These methods seek pairing costs that more accurately reflect the cost of a given pairing in operations. After these costs are found, a set partitioning model is solved.

One approach is to find a linear approximation of the expected crew cost. For any crew schedule C, let $\overline{c}(C)$ be its expected crew cost in operations. If pairing costs χ_q exist such that

$$\overline{c}(C) \; = \; \sum_{q \in C} \chi_q, \tag{19}$$

for all crew schedules C, then an optimal solution to the stochastic crew scheduling problem can be found by solving the set-partitioning problem using such pairing costs.

In general, such costs do not exist. Schaefer 2000 gives an example where costs χ_q such that (19) is satisfied for all crew schedules C do not exist. This example exploits the fact that the expected operational cost of a pairing may depend on the other pairings in the crew schedule. We seek "good" costs of this type in order to obtain an approximate solution.

4.1 The Expected Cost of a Pairing

Pairings interact when the cost of a given pairing depends on other pairings in the schedule. Interactions occur because pairings share resources such as planes, gates, flight attendants, and passengers. However, the only ways in which pairings may interact directly in our model is through shared planes and recovery. We make the following assumptions to find pairing costs that satisfy (19).

Assumption A1: The planes are always available.

Assumption A2: The recovery method is push-back, so that the departure of each flight is delayed until the crew is available and the scheduled departure time has passed.

These assumptions do not hold in practice. Crews often must change planes, and airlines often use recovery policies other than push-back. We ran an experiment to check the impact of Assumption A1 within our model of airline operations by considering a set of 136 crew schedules and simulating each for 10,000 days of airline operations in SimAir. For experiment A we used the planned routing for this fleet. If the plane was delayed from a previous flight, it may not be available for its next flight, even if the crew is available. For experiment B we used Assumption A1, so that the planes were always available. It appears from these experiments that Assumption A1 is reasonable for measuring FTC. Between experiment A and B the average operational FTC decreased by an average of 0.0986. However, the variance was very small: 9.76×10^{-5}. This indicates that although Assumption A1 does not hold in practice, the reduction in FTC is nearly constant across crew schedules. Assumption A2 does not capture all recovery options at a hub. At hubs airlines have many more options, and do allow crews to fly pairings other than the ones to which they were assigned in planning. At spokes there may be no reserve crews available or crews available for swaps, so Assumption A2 may reflect the only option available to airlines.

Under Assumptions A1 and A2 for any pairing q we define its *expected operational cost in isolation*, $\overline{c}_q$, to be the expected operational cost of a crew schedule under the push-back recovery heuristic that consists only of pairing q.

Theorem 1 *In the model of airline operations given in Section 3, under Assumptions A1 and A2 pairing costs* $\overline{c}_q$ *satisfy (19).*

We give a sketch of the full proof given in Schaefer 2000. The proof considers any pairing q for two different cases. The first assumes that the schedule consists only of pairing q and the other assumes Assumptions A1 and A2. The proof shows by induction that for any sample path of delays the operational departure and arrival time of every leg in q is the same in both cases. This implies that the pairing costs are the same in both cases, and hence the operational cost of both crew schedules is the same. Since this holds for any sample path, Theorem 1 holds.

4.2 Calculating $\overline{c}$ Pairing Costs

Since there is unlikely to be a simple formula for $\overline{c}_q$, we use a Monte Carlo simulation to estimate $\overline{c}_q$. This simulation is similar to SimAir, except that while SimAir simulates an entire fleet, this method simulates one pairing for a number of days. Let *MAXSAMPLE* and *MINSAMPLE* be two positive integers, with *MINSAMPLE* < *MAXSAMPLE*. The number of days in the sample varies by pairing. The algorithm simulates at least *MINSAMPLE* days of operations. It stops sampling when one of two criteria is satisfied:

1. The estimated pairing cost confidence interval width is less than a preset limit.

2. The algorithm has simulated *MAXSAMPLE* days of operations.

Consider the i^{th} day of operations in the simulation of pairing q. We assume that pairing q starts on time. Let pairing q consist of duties $d_1, \ldots, d_k$ where d_i consists of legs $l_{i,1}, \ldots l_{i,m(i)}$. For any leg l in pairing q, let θ_l be the minimum amount of time the crew needs after leg l for sit or rest. A pairing violates planning rules in operations when it flies more than 9 hours in any 24 hour period which does not contain a rest of at least 9 hours. We assume that the controller chooses to call a reserve crew after such a violation.

Initially, the algorithm sets *totalcost* and *samplecount* to 0. The crew's ready time for its next flight, *ctime* is initialized to $\underline{dep}(l_{1,1})$, the scheduled departure time of the first leg in the pairing. For each sample, the algorithm then loops through every duty d in the pairing q and every leg l in duty d. Prior to the departure of each leg l, the algorithm takes an observation from the ground time distribution, and then sets the operational departure time of leg l to the maximum of the crew's ready time and the scheduled departure time plus this observation. The algorithm takes an observation of block time error, and sets the arrival time of the leg to be the operational departure time plus the planned block time plus this error observation. If the crew has violated 8-in-24 planning, a reserve crew is called to fly the remainder of the pairing. The crew's next ready time is the arrival time of the current leg plus the minimum turn or rest time. At the end of each duty the operational duty cost is calculated and the subsequent rest is found. At the end of each pairing the operational pairing cost is calculated, *totalcost* is increased by this amount, and *samplecount* is incremented by 1. If the sample size is larger than *MAXSAMPLE* the algorithm terminates. If the sample size is at least *MINSAMPLE* and the confidence interval is sufficiently small the algorithm terminates. Upon termination, the algorithm returns *totalcost* divided by *samplecount*.

Algorithm 1

Initialize totalcost $= 0$, *samplecount* $= 0$

Initialize terminate to FALSE

while *terminate is FALSE* ***do***

Set $\widetilde{nctime} = \underline{dep}(l_{1,1})$

for $1 \leq i \leq k$ ***do***

for $1 \leq j \leq m(i)$ ***do***

Let ξ *be an observation from the ground time distribution*

Set $\widetilde{dep}(l_{i,j}) = \max(\underline{dep}(l_{i,j}), \widetilde{nctime}) + \xi$

Let ω *be an observation from the block time error distribution*

Set $\widetilde{arr}(l_{i,j}) = \widetilde{dep}(l_{i,j}) + \underline{block}(l_{i,j}) + \omega$

if *the crew illegal under 8-in-24 planning rules* ***then***

Call a reserve crew to fly the remainder of the pairing

end if

if *the crew will become illegal under 8-in-24 planning rules if it flies the remainder of its pairing as planned* ***then***

Call a reserve crew to fly the remainder of the pairing

end if

if $j < m(i)$ ***then***

Set $\widetilde{nctime} = \widetilde{arr}(l_{i,j}) + \theta_{l_{i,j}}$

else

if $i < k$ ***then***

if *compensatory rest is required* ***then***

Let rest be the amount of rest given by the recovery policy. This rest must be at least as long as the compensatory rest

else

Let rest be the amount of rest given by the recovery policy. This rest must be at least as long as the minimum rest

end if

Set $\widetilde{nctime} = \widetilde{arr}(l_{i,j}) + rest$

Calculate the operational duty cost

else

Calculate the operational pairing cost.

if *samplecount* $\geq$ *MINSAMPLE and totalcost/samplecount is within a given confidence interval* ***then***

Set terminate to TRUE

end if

if *samplecount = MAXSAMPLE* ***then***

Set terminate to TRUE

end if

end if

end if

end for

end for

Increment totalcost by the current pairing cost and increment samplecount by 1

end while

Return $\overline{c}_q = totalcost/samplecount$

In our experiments, *MINSAMPLE* was set to 50 and *MAXSAMPLE* was 500. We used a 99% confidence level for the termination criterion.

4.3 A Lower Bound on the Expected Cost of an Optimal Crew Schedule

In this section we give a method that finds a lower bound on the optimal objective function value for the problem of crew scheduling under uncertainty if no 8-in-24 regulations are considered in operations. Let q be any pairing, and let $\overline{o}_q$ be the expected cost of pairing q as calculated by Algorithm 1, except that operational 8-in-24 regulations are not considered. For any crew schedule C define

$$\overline{o}(C) \;=\; \sum_{q \in C} \overline{o}_q, \tag{20}$$

and let $\hat{o}(C)$ be the expected cost of crew schedule C ignoring planning 8-in-24 rules as measured by SimAir.

Theorem 2 *Under a push-back recovery heuristic and ignoring operational 8-in-24 violations, for any crew schedule C,*

$$\overline{o}(C) \leq \hat{o}(C). \tag{21}$$

The proof is similar to that of Theorem 1, and appears in Schaefer 2000. Notice that the difference between the two cases is the interaction among pairings. The proof considers a sample path of delays and shows that by ignoring the interactions among pairings the operational crew schedule costs are no greater. The reason for this is that when pairings interact one crew may need to wait for a plane before flying a given leg. While the operational flying times are not affected, operational duty elapsed time and TAFB can increase when pairings interact. Since this is true for any sample path of delays Theorem 2 holds.

If 8-in-24 planning violations are considered in operations, it is conceivable that by waiting for a plane to arrive a crew will avoid an 8-in-24 violation, and therefore Theorem 2 may no longer hold. Also, it is possible that observing 8-in-24 violations of planning rules in operations could actually lower the operational cost of a pairing. Consider a pairing q where the sum of the duty costs is the dominant factor in its planned cost. Let d be a duty flown by q where elapsed time is the dominant factor in the planned cost of d. Suppose that compensatory rest were given prior to the departure of duty d so that duty d begins late. Suppose that a large ground delay is observed at the end of duty d right after the crew has changed planes, so that the operational elapsed time of duty d is much longer than the planned elapsed time of the duty. By starting duty d late, observing 8-in-24 violations of planning rules will actually reduce the operational cost of duty d, since the operational elapsed time is lower. It is possible that the operational cost of pairing q will also be lower.

Let $\overline{C}$ be a solution to

$$\{\min\ \overline{o}x\ :\ Ax\ =\ 1,\ x \text{ binary}\}. \tag{22}$$

Corollary 3 *Let C^* be an optimal crew schedule in the sense that no crew schedule has a lower expected cost as measured by SimAir under push-back ignoring 8-in-24 planning rules in operations. Then*

$$\overline{o}(\overline{C})\ \leq\ \hat{o}(C^*). \tag{23}$$

Proof:

Immediate from Theorem 2. □

4.4 A Penalty Method

Certain attributes of a pairing may lead it to perform poorly in operations. A pairing may be close to operational limits. A pairing may contain a duty that is close to operational limits on flying time or elapsed time. A pairing containing such a duty may become illegal if it is subjected to delays. Violating such rules may result in illegal pairings.

A pairing may remain legal in operations, but still perform poorly in operations. A pairing with short sits may not be able to absorb delays without undertaking some recovery action. Similarly, a pairing with short rests may not be able to start the subsequent duty without delay.

One approach is to penalize certain attributes of pairings that may lead to poor performance in operations. In this section we describe a method of finding the optimal crew schedule for a given set of penalties. We then give a local search method for finding a best set of penalties. The hope is that the crew schedule resulting from the best set of penalties will perform well in operations.

4.4.1 Formulation

Consider any k attributes of pairings, where k is a positive integer. Examples of the attributes we consider include the number of sit minutes when the crew changes planes, the number of minutes of rest, the number of minutes of elapsed duty time and the number of minutes of duty flying time. Let the penalty space, Y, be a subset of $\mathbb{R}^{2k}$. For any attribute $1 \leq i \leq k$, let $n(i,q)$ be the number of attributes of type i that pairing q has. For instance, if attribute i is the number of minutes for a sit where the crew changes planes, $n(i,q)$ is the number of sits where the crew changes planes. Let $a_q^{(i,j)}$ be the value of attribute i for pairing q, where $1 \leq j \leq n(i,q)$ and $1 \leq i \leq k$. For instance, if pairing q has 5 sits where the crew changes planes, with lengths 55, 72, 63, 58 and 67 respectively, then $n(i,q) = 5$, $a_q^{(i,1)} = 55$, $a_q^{(i,2)} = 72$, $a_q^{(i,3)} = 63$, $a_q^{(i,4)} = 58$, and $a_q^{(i,5)} = 67$.

For any attribute i, $1 \leq i \leq k$, let ϑ_i be the largest or smallest amount that is acceptable for that particular attribute. For the attribute corresponding to the elapsed time of a duty, ϑ is the maximum elapsed time permitted for a duty, but for sits, ϑ is the length of the minimum legal sit. For instance, the sit attribute identified above may have a ϑ value of 45 minutes; sits shorter than 45 minutes are not permitted. Let α_i and γ_i be positive real numbers. We interpret α_i as the maximum penalty for factor i and γ_i as the slope of the penalty function. Consider any penalty combination $(\alpha, \gamma) = (\alpha_1, \ldots, \alpha_k, \gamma_1, \ldots, \gamma_k) \in Y$. Then the function $f_i(\cdot)$ is defined as

$$f^i(\alpha_i, \gamma_i, q) = \sum_{j=1}^{n(i,q)} \max\Big(\alpha_i - \gamma_i(a_q^{(i,j)} - \vartheta_i), 0\Big). \tag{24}$$

For any attribute, as $a_q^{(i,j)}$ approaches ϑ_i, the resulting schedule may be more likely to have disruptions due to that attribute. For example, as the sit time decreases, the more likely it will be that a flight must be delayed because the crew is unavailable. Whenever $a_q^{(i,j)}$ exceeds $\frac{\alpha_i}{\gamma_i} + \vartheta_i$, the penalized amount is 0 for that particular j.

For any $(\alpha, \gamma) \in Y$, let $x^*(\alpha, \gamma)$ be the optimal solution to the deterministic crew scheduling problem with pairing costs

$$c_q \;=\; \underline{c}_q \;+\; \sum_{i=1}^{k} f^i(\alpha_i, \gamma_i, q).$$

Let $s(\alpha, \gamma)$ be the expected cost of crew schedule $x^*(\alpha, \gamma)$ as estimated through SimAir. To determine $s(\alpha, \gamma)$, crew schedule $x^*(\alpha, \gamma)$ is simulated by SimAir. We would like to solve

$$\min_{\alpha, \gamma \in Y} s(\alpha, \gamma). \tag{25}$$

Unfortunately, this problem is very difficult to solve. In general, $s(\alpha, \gamma)$ is not continuous, and for any given $(\alpha, \gamma) \in Y$, finding $s(\alpha, \gamma)$ is quite expensive in that it requires solving a deterministic crew scheduling problem. The computational results given in this section demonstrate that $s(\alpha, \gamma)$ is neither convex nor concave. It seems unlikely that a global optimum to problem (25) can be found. We propose a local search of the penalty space to find crew schedules that perform well in operations. The goal is to find an approximate solution $(\hat{\alpha}, \hat{\gamma})$ to problem (25).

While this methodology could be extended to find a crew schedule that performs well for other criteria, such as on-time performance, we will evaluate the quality of a crew schedule based on its expected crew cost in operations as estimated by SimAir.

4.4.2 Local Search of Penalty Space

For our experiments we used the four attributes listed below.

Attribute 1 The number of minutes of scheduled sit when the crew is scheduled to change planes, hereafter referred to as *swap time*. The parameter ϑ_1 is set to 45 minutes.

Attribute 2 The number of minutes of scheduled rest time between duties. The parameter ϑ_2 is set to 615 minutes, or 10 hours, 15 minutes.

Attribute 3 The number of minutes of flying in a duty. The parameter ϑ_3 is set to 480 minutes, or 8 hours.

Attribute 4 The number of minutes of elapsed time in a duty. The parameter ϑ_4 is set to 810 minutes, or 13 hours and 30 minutes.

The local search procedure starts with the deterministic or planning solution, with all penalty levels set to 0. It then varies attributes 1 through 4 sequentially. The algorithm maintains an incumbent solution that has the lowest expected cost in operations, as measured by a Monte Carlo estimate from SimAir. If an improved schedule is found, that is, a schedule with better expected performance than the incumbent solution, the incumbent solution is replaced by the improved solution. After the fourth factor has been considered, the first factor is again varied to see if any improvement is possible. If no improvement has been found, the algorithm terminates and the incumbent solution is returned. Otherwise, if an improved solution is found, we reexamine the other factors.

4.5 Finding the Crew Schedules

Once an objective function is determined, we solve the set partitioning problem using an algorithm developed by Klabjan et al. 1999c. A large set of pairings is generated in parallel. The algorithm divides the starting duties among the processors. The pairings are found by depth-first search through a duty time-line network. The algorithm can either enumerate all legal pairings for a fleet, or a random subset of the legal pairings. For larger fleets the large number of legal pairings makes enumeration impractical. A random subset of duties is found, and then a random depth-first search algorithm enumerates a subset of the pairings containing the randomly-generated duties. For more details about pairing generation, see Klabjan and Schwan 1999.

The algorithm solves an LP relaxation over the generated pairings using the parallel primal-dual simplex algorithm developed by Klabjan et al. 1999b. This algorithm is able to solve linear programs with millions of columns. The algorithm solves linear programs over a subset of columns on the various processors and it combines the dual solutions to obtain a dual-feasible solution. The algorithm removes all nonbasic columns and randomly generates a new set of columns. It repeats this process until termination criteria are met. A smaller set of pairings is chosen based on reduced cost and a random selection heuristic. Finally, the set partitioning problem is solved over this smaller set of pairings.

5 Computational Results

We considered three daily fleets provided to us by a major domestic carrier. We refer to these fleets as F1, F2, and F3. For Fleet $i = 1, 2, 3$, let $\underline{C}_{Fi}$ be the deterministic crew schedule, let $\overline{C}_{Fi}$ be the crew schedule found by the expected cost method given in Algorithm 1, and let $\hat{C}_{Fi}$ be the best crew schedule found by the penalty method.

5.1 Computational Results for Fleet F1

Fleet F1 has about 120 daily legs. Let The results for Fleet F1 is given in Table 1. The lower bound on C^* for Fleet F1 found by ignoring operational 8-in-24 rules is 4.10. Thus, crew schedule $\overline{C}_{F1}$ has an expected FTC that is nearly equidistant between the performance of $\underline{C}_{F1}$ and the lower bound.

Table 1: Summary for Fleet F1.

Schedule	Planned FTC	Operational FTC	On-time percentage
$\underline{C}_{F1}$	2.51	4.31	73.9
$\overline{C}_{F1}$	2.64	4.22	73.8
$\hat{C}_{F1}$	2.51	4.29	73.9

5.2 Computational Results for Fleet F2

Fleet $F2$ has about 150 daily legs. The performances in SimAir are summarized in Table 2. The lower bound found by ignoring operational 8-in-24 rules for Fleet $F2$ is 8.40. The gap between the performance of crew schedule $\overline{C}_{F2}$ and the lower bound is approximately 20% smaller than the gap between $\underline{C}_{F2}$ and the lower bound. The penalty method was not successful for this fleet and could not improve upon $\underline{C}_{F2}$ nor $\overline{C}_{F2}$ in operations.

Table 2: Summary for Fleet F2.

Schedule	Planned FTC	Operational FTC	On-time percentage
$\underline{C}_{F2}$	3.79	9.11	79.4
$\overline{C}_{F2}$	3.92	8.97	79.2

5.3 Computational Results for Fleet F3

Fleet F3 has over 340 daily legs. The computational results for Fleet F3 are summarized in Table 3. The lower bound on the optimal solution for this fleet is 5.51. Thus, the gap between crew schedule $\overline{C}_{F3}$ and the lower bound is 60% as large as the gap between $\underline{C}_{F3}$ and the lower bound.

Table 3: Summary for Fleet F3.

Schedule	Planned FTC	Operational FTC	On-time percentage
$\underline{C}_{F3}$	2.69	5.82	72.6
$\overline{C}_{F3}$	2.91	5.69	72.6
$\hat{C}_{F3}$	2.74	5.75	72.5

For each of the three fleets the crew schedule found by Algorithm 1 performed better than the crew schedule found using deterministic methodology, and has also provided substantially better results than the penalty method. Relative to the lower bounds established by ignoring operational 8-in-24 rules Algorithm 1 performed noticeably better than the deterministic method. The difference between the cost of the schedules found by Algorithm 1 and the lower bound was more than 50% smaller than the difference between the cost of the deterministic schedules and the lower bound for two of the fleets.

5.4 An Analysis of the Crew Schedules

We consider the deterministic crew schedules and the schedules found by Algorithm 1 for each of the the three fleets. For each crew schedule C found for each of the fleets we express the difference between $\sum_{q \in C} \overline{c}_q$ and crew schedule C's planned cost as a percentage of the total increase between the operational and planned cost of crew schedule C. Mathematically, this is expressed as

$$\frac{\sum_{q \in C} \overline{c}_q \ - \ \underline{c}(C)}{\overline{c}(C) \ - \ \underline{c}(C)} \ \cdot \ 100. \tag{26}$$

The results are displayed in Table 4. The consistently large percentages indicate that possible interactions among pairings have a small impact on the total difference between a crew schedule's operational and planned cost. This provides further empirical evidence that Assumption A1 appears to be reasonable in our model of airline operations.

Table 4: The Effect of Pairing Interactions.

Crew Schedule	Cost Increase Explained by $\overline{c}$ Costs
$\underline{C}_{F1}$	94 %
$\overline{C}_{F1}$	92 %
$\underline{C}_{F2}$	90 %
$\overline{C}_{F2}$	89 %
$\underline{C}_{F3}$	96 %
$\overline{C}_{F3}$	94 %

We now analyze the schedules to determine how often the crew follows the planes according to the actual routing, and how many pairings and duties are in both $\underline{C}_{Fi}$ and $\overline{C}_{Fi}$ for $i = 1, 2, 3$. We give the factor dominating the pairing costs, as well as the largest deterministic FTC for each fleet. The results are given in Table 5.

Table 5: A comparison of crew schedules $\underline{C}$ and $\overline{C}$ for the three fleets.

Crew Schedule	Matching pairings	Matching duties	Crew follows plane	$\sum$ duty is max	TAFB is max	MG is max	Max FTC
$\underline{C}_{F1}$	12/25	63/72	21/47	15/25	10/25	0/25	12.7
$\overline{C}_{F1}$	12/24	63/73	22/46	13/24	11/24	0/24	5.2
$\underline{C}_{F2}$	7/14	34/41	47/108	4/14	5/14	5/14	30.4
$\overline{C}_{F2}$	7/15	34/43	46/106	4/15	6/15	5/15	22.7
$\underline{C}_{F3}$	17/42	83/124	68/218	20/42	20/42	2/42	16.2
$\overline{C}_{F3}$	17/41	83/124	75/218	17/41	19/41	5/41	14.5

For these three fleets, fewer pairings determined by Algorithm 1 had the sum of the duty costs as the dominant factor in their costs. Intuitively, this makes sense; TAFB depends largely on the end of the pairing, so in most cases it has a smaller variance than the sum of the duty costs. By choosing more pairings where TAFB is the largest factor in planning, Algorithm 1 is able to choose from a richer set of pairings. By doing this, it is able to avoid pairings with large deterministic FTCs. It is able to recognize that even though a pairing may have TAFB or minimum guarantee dominate in planning, it does not necessarily mean that this will remain the case in operations. This may be why Algorithm 1 appears to perform better in operations; it is able to consider a wider range of pairings that are likely to be paid for flying time in operations, rather than the smaller set that is paid for flying time in planning.

Several patterns emerged across all three fleets. Pairings with 0 planned FTCs had larger differences between their planned and operational FTCs than pairings with positive planned FTCs. This has several implications. First, a pairing with a small planned FTC may be equally desirable in operations as a pairing with a 0 planned FTC. Second, given two crew schedules with equal total planned FTC, it may be preferable to choose a crew schedule with many pairings with small planned FTCs over a crew schedule with many 0 planned FTC pairings and a few pairings with large planned FTCs. For all three fleets, the expected cost schedule had fewer 0 planned FTC pairings than the deterministic schedule. Because the expected cost algorithm views more pairings as acceptable, it is able to avoid using pairings with large planned FTCs. Given two crew schedules with equal planned FTCs, having many pairings with positive planned FTCs appears to be more desirable than having a few. Pairings with small planned FTCs still may have the sum of duty costs dominate in operations; this is unlikely for pairings with large planned FTCs. It also appears that the cost due to interaction between pairings is insignificant compared to the cost arising from each pairing considered in isolation. Isolating pairings allows us to use the standard set-partitioning model for solving these crew scheduling problems. This is a significant finding, since explicitly considering interactions between pairings would make solving crew scheduling problems even more difficult.

Using the pairing costs $\bar{c}$ found by Algorithm 1 in the standard set-partitioning crew scheduling model results in crew schedules that perform significantly better than deterministic crew schedules in the model of airline operations used by SimAir. A significant reduction in operational crew costs may be found by considering each pairing in isolation and then using its expected operational cost in the objective function of the crew scheduling problem. One insight provided by these results is that pairings with small planned FTCs may in fact perform well in operations, since pairings with 0 planned FTCs appear to have the largest difference between operational and planned FTCs. Algorithm 1 recognizes this, and hence it has fewer 0 planned FTC pairings than the deterministic crew schedules.

6 Acknowledgements

This research was sponsored by grants from United Airlines Research and Development and National Science Foundation grant DMI-9700285. The authors would like to thank Srini Ramaswamy of United Airlines, Eric Gelman of American Airlines, and Dave Goldsman and Jay Rosenberger of Georgia Tech for their suggestions and insight. The computational experiments were made using hardware donated by the Intel Corporation and optimization software provided by ILOG.

References

R. Anbil, F. Barahona, L. Ladyani, R. Rushmeier, and J. Snowden, "IBM Makes Advances in Airline Optimization", Technical Report RC21465(96887)04MAY1999, IBM Research, Yorktown Heights, NY, 1999.

R. Anbil, E. Gelman, B. Patty, and R. Tanga, "Recent Advances in Crew-Pairing Optimization at American Airlines", *Interfaces* 21, 62–74 (1991).

R. Anbil, L. Hatay, C. Barnhart, E. L. Johnson, and V. S. Ramakrishnan, "Crew-Pairing Optimizaton at American Airlines Decision Technologies," in *Optimization in Industry*, T. A. Cirani and R. C. Leachman (eds), 31–36, John Wiley and Sons, New York, 1992a.

R. Anbil, R. Tanga, and E. L. Johnson, "A Global Approach to Crew-Pairing Optimization", *IBM System Journal* 31, 71–78 (1992b).

E. Andersson, E. Housos, N. Kohl, and D. Wedelin, "Crew Pairing Optimization," in *Operations Research in the Airline Industry*, G. Yu (ed.), 228–258, Kluwer Academic Publishing, Amsterdam, 1997.

Anonymous, "A Jam at 32,000 Feet", *The Economist*, February 5, 2000.

C. Barnhart, E. L. Johnson, L. Hatay, and R. Anbil, "A Column-Generation Technique for the Long-Haul Crew-Assignment Problem," in *Optimization in Industry 2*, T. A. Cirani and R. C. Leachman (eds), 31–36, John Wiley and Sons, New York, 1994.

J. F. Benders, "Partitioning Procedures for Solving Mixed Variables Programming Problems", *Numerische Mathematics* 4, 238–252 (1962).

BTS, "On-Time Statistics", available from http://www.bts.gov/ntda/oai/desc.html (1998).

B. Caldwell, "Delta Powers up", *IT Management* 1, 613 (1997).

H. D. Chu, E. Gelman, and E. L. Johnson, "Solving Large Scale Crew Scheduling Problems", *European Journal of Operations Research* 97, 260–268 (1997).

FAA, "Federal Aviation Regulations", available from http://www.faa.gov/avr/AFS/FARS/far-121.txt (1999).

I. Gershkoff, "Optimizing Flight Crew Schedules", *Interfaces* 19, 29–43 (1989).

G. W. Graves, R. D. McBride, I. Gershkoff, D. Anderson, and D. Mahidhara, "Flight Crew Scheduling", *Management Science* 39, 736–745 (1993).

K. L. Hoffman and M. Padberg, "Solving Airline Crew Scheduling Problems by Branch-and-Cut, *Management Science* 39, 657–682 (1993).

D. Klabjan, E. L. Johnson, and G. L. Nemhauser, "Airline Crew Scheduling with Regularity", Technical Report TLI/LEC-99-11, Georgia Institute of Technology, Atlanta, GA, 1999a.

D. Klabjan, E. L. Johnson, and G. L. Nemhauser, "A Parallel Primal-Dual Simplex Algorithm", Technical Report TLI/LEC-99-08, Georgia Institute of Technology, Atlanta, GA, 1999b.

D. Klabjan, E. L. Johnson, and G. L. Nemhauser, "Solving Large Airline Crew Scheduling Problems: Random Pairing Generation and Strong Branching", Technical Report TLI/LEC-99-12, Georgia Institute of Technology, Atlanta, GA, 1999c.

D. Klabjan and K. Schwan, "Airline Crew Pairing Generation in Parallel", Technical Report TLI/LEC-99-09, Georgia Institute of Technology, Atlanta, GA, 1999.

S. Lavoie, M. Minoux, and E. Odier, "A New Approach for Crew Pairing Problems by Column Generation with an Application to Air Transportation", *European Journal of Operational Research*, 35, 45–58 (1988).

L. Lettovský, Airline Operations Recovery: An Optimization Approach, Ph.D. thesis, Georgia Institute of Technology, Atlanta, GA (1997).

A. W. Mathews, "Busier Skies, More Stress on Systems Foreast by FAA for the Next Decade", *The Wall Street Journal*, March 7, 2000 A3.

R. O'Dell, "Airline Delays Skyrocket 58% from 1995 to '99, Report Says", *The Atlanta Constitution*, July 26, 2000 E1.

D. Phillips and N. Irwin, "Flying into a Storm of Delays", *The Washington Post*, July 17, 2000, A01.

J. M. Rosenberger, A. J. Schaefer, D. Goldsman, E. L. Johnson, A. J. Kleywegt, and G. L. Nemhauser, "SimAir: A Stochastic Model of Airline Operations", in *Proceedings of the 2000 Winter Simulation Conference*, J. Joines, R. R. Barton, P. Fishwick and K. Kang (eds), 1118–1122, Institute of Electrical and Electronics Engineers, Piscataway, NJ, 2000a.

J. M. Rosenberger, A. J. Schaefer, D. Goldsman, E. L. Johnson, A. J. Kleywegt, and G. L. Nemhauser, "A Stochastic Model of Airline Operations", Technical Report TLI/LEC-00-06, Georgia Institute of Technology, Atlanta, GA, 2000b.

A. J. Schaefer, Airline Crew Scheduling under Uncertainty, Ph.D. thesis, Georgia Institute of Technology, Atlanta, GA (2000).

J. Snowden, R. Anbil, and G. Pangborn, "The Airline Crew Scheduling Problem: Dual Simplex, Volume, and Volume/SPRINT Solutions", Technical Report RC21642(01/14/00), IBM Research, Yorktown Heights, NY, 2000.

P. H. Vance, A. Atamturk, C. Barnhart, E. Gelman, E. L. Johnson, A. Krishna, D. Madidhara, G. L. Nemhauser, and R. Rebello, "A Heuristic Branch-and-Price Approach for the Airline Crew Pairing Problem", Technical Report TLI/LEC-97-06, Georgia Institute of Technology, Atlanta, GA, 1997.

J. W. Yen, A Stochastic Programming Formulation of the Stochastic Crew Scheduling Problem, Ph.D. thesis, University of Michigan, Ann Arbor, MI (2000).