# A Stochastic Model of Airline Operations

Jay M. Rosenberger
*Operations Research and Decision Support Department, American Airlines, Dallas-Fort Worth Airport, Texas 75261-9616*

Andrew J. Schaefer
*Department of Industrial Engineering, School of Engineering, University of Pittsburgh, 1048 Benedum Hall, Pittsburgh, Pennsylvania 15261*

David Goldsman • Ellis L. Johnson • Anton J. Kleywegt
George L. Nemhauser
*School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332-0205*

We present a stochastic model of the daily operations of an airline. Its primary purpose is to evaluate plans, such as crew schedules, as well as recovery policies in a random environment. We describe the structure of the stochastic model, sources of disruptions, recovery policies, and performance measures. Then, we describe SimAir—our simulation implementation of the stochastic model, and we give computational results. Finally, we give future directions for the study of airline recovery policies and planning under uncertainty.

Airline operations are subject to many uncertainties. Anecdotal evidence suggests that major domestic airline carriers almost never experience a day without disruptions. A *disruption* is an event that prohibits an airline from operating as scheduled. Recent newspaper articles on the causes and implications of disruptions include Adams (2000), Dobbyn (2000), Phillips and Irwin (2000), and Reuters (2000), and some facts from them are given below. Average daily flight delays increased 20% from 1998 to 1999, and customer complaints increased by 130%. In June 1999, the airlines began a "Customer First" campaign to avoid additional government regulation. Unfortunately, airline performance has not improved. In May 2000, there were four days with over 2,000 air traffic delays each. In June 2000, flight delays increased 16.5% from June 1999.

Several factors in airline systems account for the current frequency of disruptions. There are more passengers, and planes are more crowded than at any time in history. Air traffic controllers are struggling to accommodate the number of flights the airlines want to offer. Mechanical failures can disrupt the planned schedule. On June 10, 2000, 71 flights were cancelled due to mechanical problems by one major domestic carrier. Weather accounts for nearly 75% of all delays, and in June 2000, weather accounted for as many as 79%.

Even though disruptions occur frequently, current airline planning models do not explicitly consider disruptions in operations; therefore, an airline's actual performance can be significantly different from the planned performance. Traditional airline planning models measure the quality of a plan assuming that every flight takes off and lands as scheduled. Because this optimistic scenario rarely occurs, a better measure of the quality of a plan is its performance in *operations*, when the plan is executed. It is not easy to determine the performance of a plan in operations a priori due to unknown future disruptions.

0041-1655/02/3604/0357$05.00
1526-5447 electronic ISSN

One challenge in evaluating the future performance of a plan in operations is to foresee how recovery will take place. *Recovery* is the way in which an airline reacts to disruptions. Flights may be delayed or cancelled, and pilots or planes may be rescheduled. Different recovery policies will give different performance results.

We present a stochastic model of the daily operations of an airline. Its primary purpose is to evaluate the performance of plans and recovery policies in operations. By building a better understanding of airline operations and how plans and recovery policies affect them, the airlines could improve their on-time performance, reduce their cost in operations, and increase customer satisfaction. We describe optimization techniques for airline planning and airline operations in §1. Section 2 describes the structure of the model and sources of airline disruptions, and §3 considers components of recovery policies. Section 4 discusses performance measures for evaluation. In §5 we describe SimAir—our simulation implementation of the stochastic model. Section 6 presents examples, and §7 suggests directions for the further study of airline planning and recovery under uncertainty.

There are several simulation implementations of stochastic airline models. Yang et al. (1991) developed an airline simulation for aircraft reliability. Their implementation does not explicitly consider crews or passengers, and their recovery policy for flight cancellations is simpler than ours. Haeme et al. (1988) developed an airline simulation that considers crews and passengers to assist in schedule development. Their implementation uses a recovery policy similar to the default recovery policy for SimAir, but it does not support more sophisticated recoveries. Yau (1991) described a simulation within an airline planning decision support system. The focus of his decision support system is for short-term airline planning; it does not include crew recovery and long-term scheduling.

## 1. Overview of Airline Optimization

### 1.1. Optimization in Airline Planning

There are multiple stages in airline planning. In practice, the stages given below are solved sequentially. Yu (1997) provides more detail about optimization in airline planning and operations. We discuss only the stages relevant to our stochastic model.

**1.1.1. Flight Schedule.** The first stage of the airline planning process is to develop the flight schedule. A *station* is an airport that an airline serves. A *leg* consists of an origin station, a destination station, a departure time, and an arrival time. The *block time* of a leg is the length of time from the moment the plane leaves the gate at the origin station until the plane arrives at the gate of the destination station.

Flight scheduling determines the origin, the destination, the departure time, and the arrival time of each leg. Typically, this takes place at least 3 to 12 months in advance of the scheduled legs and is driven largely by market considerations. A flight schedule is subject to minor changes until the departure of each leg.

A *hub-and-spoke* network refers to the structure of an airline flight schedule in which a large percentage, as many as 98%, of the legs go into or out of a small subset of stations called *hubs*. *Spokes* have limited activity. Most major domestic airlines use hub-and-spoke networks, typically with five or six hub stations. Hub-and-spoke networks allow passengers to fly from an origin to a destination with very few intermediate stops. A *passenger itinerary* is a sequence of legs that typically connect at hub stations. Unfortunately, hub-and-spoke systems are highly sensitive to disruptions at hubs. A disruption at a hub station can prevent many passengers from flying their original itineraries. Passengers on a *misconnecting itinerary* do not have sufficient time to make one or more of their connections when changing planes. At least one leg is cancelled in a *cancelled itinerary*. Passengers on misconnecting and cancelled itineraries are rerouted on other legs.

**1.1.2. Fleet Assignment.** Large domestic airline carriers usually have more than one type of aircraft. A *fleet* is a set of planes of the same type. Two planes of the same fleet type can have different passenger capacities. For example, a Boeing 737-300 carries 138 passengers, while a 737-200 has 111 seats (Olympic Airways 2001). After the initial schedule is set, each leg is assigned to a fleet with a specific passenger capacity. Fleet assignment problems are solved using an integer multicommodity flow model. Hane et al. (1995), Gu et al. (1994), Barnhart et al. (1998), Boland et al. (2000), and Kniker (1998) describe the fleet assignment model in detail.

**1.1.3. Aircraft Rotation.** After the legs are assigned a fleet type, airline planners create an aircraft rotation. An *aircraft rotation* is a sequence of legs flown by the planes within the same fleet. Many airlines require the planes in the same fleet to fly the same rotation. A rotation generally takes many days to fly. A *daily route* is a subsequence of a rotation occurring within the same day. The rotation must comply with certain maintenance restrictions requiring periodic plane service called *scheduled maintenance*. Clarke et al. (1997) describe an aircraft rotation algorithm.

In addition to the scheduled maintenance, when planes experience mechanical problems in operations they receive *unscheduled maintenance*. A rotation that considers the possibility of an unscheduled maintenance problem could be less sensitive to aircraft mechanical problems. Yang et al. (1991) develop a simulation that models unscheduled maintenance problems.

**1.1.4. Crew Scheduling.** The crew-scheduling problem partitions the set of legs into trips that crews will fly. 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 legs that follow certain legality rules 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* is a sequence of duties that starts and ends at the same city. For any two consecutive duties $d_i$ and $d_{i+1}$ in a pairing, duty $d_i$ must finish in the same city where duty $d_{i+1}$ begins. Such 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 must adhere to certain Federal Aviation Administration (FAA) and contractual rules. For instance, one such rule requires that a crew must receive compensatory rest if the crew flies more than 8 hours within a 24-hour period. A pairing that violates this planning rule is *illegal* and cannot be included in a crew schedule. However, under extreme conditions, the FAA allows pairings that were legal in planning to violate this rule in operations (Federal Aviation Administration 1999). A partially flown pairing in operations is scheduled to violate 8-in-24 planning rules if by flying the remainder of the pairing, the crew would violate these rules. A recovery policy would typically reroute the crew to avoid such a circumstance. The *time away from base* (TAFB) of a pairing is the amount of time that elapses between the beginning of the pairing and the end of the pairing. In many instances, crews are paid based upon 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 of the duties is significantly large. We describe the details of calculating crew cost in §4.1.

A *crew schedule* is a set of legal pairings that partitions the legs of a single fleet. Crew-scheduling problems are solved by generating pairings and solving a set partitioning problem. The *daily crew-scheduling problem* is solved under the assumption that the crew schedule is repeated every day. Hoffman and Padberg (1993), Chu et al. (1997), Vance et al. (1997), Klabjan and Schwan (1999), and Klabjan et al. (1999b) describe implementations for the daily problem.

A flight schedule may vary throughout a week. For example, the flight schedule on Monday can be different from the one on Saturday. The *weekly problem* gives a crew schedule for each day of the week (Klabjan et al. 1999a).

**1.1.5. Crew Assignment.** The next planning stage assigns pilots to pairings. Crew assignment is done using a bidline or preferential model. A *bidline* is a set of pairings that a crew flies within a month. Every bidline must adhere to certain FAA and contractual rules. For example, within any seven-day period, a crew cannot be assigned to fly more than 30 hours and must be given a rest of at least 24 hours (Federal Aviation Administration 1999). We refer to these restrictions as *weekly planning rules*. A *bidline model* generates a set of bidlines, and pilots sequentially choose the bidline they prefer in order of seniority (Christou et al. 1999).

In a *preferential model*, pilots place weights on characteristics that they value in a bidline. Preferential models find the optimal set of bidlines for each pilot in order of seniority (Gamache et al. 1998).

### 1.2. Optimization in Airline Operations

Recovery is the process of reacting to a disruption. The optimal recovery decision is rarely easy to determine. The future is uncertain and changing, and canceling a leg or rerouting a crew or a plane can have profound consequences throughout the airline's system. In practice, the airlines use an *Airline Operations Control Center* (AOCC) to implement recovery. An AOCC does most recovery manually (Lettovský 1997). This fact makes airline recovery difficult to model because AOCC personnel often act upon their intuition. Most optimization research done on airline operations has been on crew recovery. These models assume all legs will fly according to their new scheduled leg times. We are unaware of any previous research on dynamic and stochastic airline recovery models.

Most airline planning models assume a daily flight schedule, although some large-scale models use a weekly flight schedule. Because a flight schedule frequently changes in operations from day to day, airline operational models distinguish between two identical legs scheduled to fly on different dates. For example, suppose the set of flights shown in Table 1 is flown every day. On Monday, the crew flying Flight 1 arrived 20 minutes late to the airport. Due to a snow storm in Minneapolis on Monday, Flights 2 and 3 were cancelled, and Flight 4 was delayed by 15 minutes. The recovery would consider the scheduled legs in Table 2. A *dated flight schedule* refers to the legs as they were flown or scheduled to be flown on a specific date. The departure and arrival times of a daily leg are the *original* departure and arrival times, whereas the departure and arrival times of a dated leg are the *scheduled* departure and arrival times. Observe from Tables 1 and 2 that the original times do not equal the scheduled times because the legs have been rescheduled.

**Table 1 Daily Flight Schedule**

| Flight | Origin | Original Departure Time | Destination | Original Arrival Time |
|---|---|---|---|---|
| 1 | San Diego | 6:45 | Burbank | 7:28 |
| 2 | Burbank | 8:25 | Minneapolis | 14:16 |
| 3 | Minneapolis | 15:00 | Burbank | 17:18 |
| 4 | Minneapolis | 19:20 | San Diego | 21:40 |
| 5 | Burbank | 18:00 | Minneapolis | 23:46 |

*Note.* Each leg is flown each day.

**1.2.1. Crew Recovery.** Lettovský et al. (2000) and Stojković et al. (1998) describe crew recovery models for airline operations. Crew recovery models use a dated crew schedule, a scenario time, and a list of disruptions. The *dated crew schedule* is a set of scheduled pairings, and the *scenario time* is the time when the crew schedule is reoptimized. The disruptions include delays and cancellations of dated legs before and after the scenario time. The delays and cancellations are inputs to the crew recovery model and are applied to the dated crew schedule. A heuristic selects a set of pairings to be reoptimized, and the crew recovery model uses the set of dated legs from the set of selected pairings to generate new pairings. Then the recovery model solves a set-partitioning problem over the set of new pairings.

**1.2.2. Integrated Recovery.** Integrated recovery models simultaneously consider crew, aircraft rout-

**Table 2 Dated Flight Schedule**

| Flight | Origin | Scheduled Departure Day | Scheduled Departure Time | Destination | Scheduled Arrival Time |
|---|---|---|---|---|---|
| 1 | San Diego | Monday | 7:05 | Burbank | 7:48 |
| 4 | Minneapolis | Monday | 19:35 | San Diego | 21:55 |
| 5 | Burbank | Monday | 18:00 | Minneapolis | 23:46 |
| 1 | San Diego | Tuesday | 6:45 | Burbank | 7:28 |
| 2 | Burbank | Tuesday | 8:25 | Minneapolis | 14:16 |
| 3 | Minneapolis | Tuesday | 15:00 | Burbank | 17:18 |
| 4 | Minneapolis | Tuesday | 19:20 | San Diego | 21:40 |
| 5 | Burbank | Tuesday | 18:00 | Minneapolis | 23:46 |

*Note.* Each dated leg is flown on a specific date.

ing, and passenger recovery problems. Integrated recovery models also determine how long to delay a leg and whether to cancel a leg. Based upon the delays and cancellations, they may reroute planes, passengers, and crews. Lettovský (1997) describes an integrated recovery model that uses a decomposition algorithm for crew recovery, aircraft recovery, 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 in the description.

## 2. Stochastic Model of Airline Operations

In this section, we present a general description of our stochastic model, which is a discrete event semi-Markov process, described in terms of states and transitions that can either be random or deterministic. A complete description of the model is in Rosenberger et al. (2000). The input of our model is an *original schedule*. This schedule includes a set of legs, a set of crews, their pairings and their bidlines, a set of planes, their rotation, a set of itineraries, their passengers, and a set of *reserve crews* that are not assigned to pairings but can be used in operations. The *current schedule* is the schedule that has been most recently updated. The *state* of the model includes:

- Information that describes how the current schedule deviates from the original schedule at the current time, so that the current schedule can be reconstructed from the state and the original schedule.
- Historical information needed to calculate performance measures and to determine whether the current schedule will violate planning rules. For example, part of the history of a pairing determines whether the crew violates 8-in-24 planning rules.
- Conditions beyond the airline's control such as weather, which can affect the airline system.

A state transition can result from either an event or a set of decisions. For example, there are departure events, arrival events, repaired plane events, weather events, and airport congestion events. After each event transition, the state may further change depending on the operational decisions defined by a recovery policy.

Because all of the legs in the model are dated, we use the term *leg* instead of *dated leg* for the remainder of this paper.

### 2.1. State

The model describes the operations of a particular airline or a particular fleet of an airline. The operations of other airlines are not modeled in the same amount of detail. Instead, the effect that other airlines have on congestion is modeled with a time-dependent semi-Markov process. The model also uses a semi-Markov process for the state of the weather.

*Restricted flow* at a station reduces the rate of departure and arrival legs. To account for the weather and airport congestion, flows at the stations that are experiencing congestion and bad weather are restricted. For each station, the state includes the following *flow components*:

- a *runway queue.* The runway queue is a sequence of legs with a service time distribution that is dependent on the congestion and the weather.
- a *taxi time.* The taxi time is the length of time a plane needs to taxi from the departure gate to the runway queue.
- an *airspace queue.* The airspace queue is a sequence of legs with a service time distribution that is dependent on the congestion and the weather.
- a *landing time.* The landing time is the length of time a plane needs to land and taxi from the beginning of the landing approach to the gate of the arrival station.

The route that a plane flies is called the plane's *flow*. When a plane arrives at a station, it must depart on its next leg from the same station. This requirement is *flow balance*. The sequence of legs flown by a pairing is the crew's flow, and the set of legs on a passenger itinerary is the passengers' flow. The current schedule must maintain flow balance for all three flows, and the state includes information about the flows. For each plane the state includes:

- The plane's currently scheduled rotation.
- Information about the plane after its most recent scheduled maintenance service. Since planes require maintenance inspections after a given amount of activity and time, the state includes:

— The time when the plane received its most recent scheduled maintenance service.

— The number of departures since the plane's most recent scheduled maintenance service.

— The number of minutes flown since the plane's most recent scheduled maintenance service.

• An indicator determining whether the plane is currently being serviced.

• The time at which the plane will be repaired if it is being serviced.

• The leg that the plane is flying or is scheduled to fly next in the assigned rotation.

For each crew the state includes:

• The number of minutes the crew has flown in the current month and the current year. The FAA and contractual rules limit crew usage within a month and within a year.

• The departure and arrival times of the legs flown by the crew in the previous seven days. These times determine whether the crew is violating or scheduled to violate weekly planning rules.

• The scheduled departure and arrival times of any legs not yet flown and the departure and arrival times of any legs already flown in the pairing to which the crew is currently assigned. The legs in the pairing determine whether the crew is violating or is scheduled to violate 8-in-24 planning rules.

• The leg that the crew is flying or is scheduled to fly next in the assigned pairing.

For each passenger itinerary the state includes:

• The current sequence of legs flown by the passengers on the itinerary. This sequence of legs determines whether the passengers are scheduled to miss at least one of their connecting flights.

• The number of passengers flying the itinerary. The number of passengers on the itinerary can determine operational decisions to reduce the number of passengers that are rerouted.

### 2.2. Events

There are many sources of delays, including overbooking, baggage loading, etc. Because our model does not explicitly consider the sources of these delays, it is unnecessary to define them individually. Instead, the model uses aggregate distributions for the *ground time*, the time duration from the moment the plane and crew are ready until the departure of a leg. A block time disruption changes the number of minutes a crew flies. A ground time disruption does not alter the number of minutes a crew flies, although it could affect the elapsed time of a duty or the time away from base of the pairing.

The model decomposes the block time of a flight according to six events. The six events are determined by the queueing network displayed in Figure 1:

• A *departure event* occurs when the plane pushes away from the gate and begins to taxi to the runway.

• A *runway event* occurs when the plane enters the runway queue of the departure station.

• A *take-off event* occurs when the plane reaches the front of the runway queue and begins its flight.

• An *airspace event* occurs when the plane enters the arrival station airspace queue.

• An *approach event* occurs when the plane reaches the front of the airspace queue and begins its landing approach.

• An *arrival event* occurs when the plane lands and taxis to the gate.

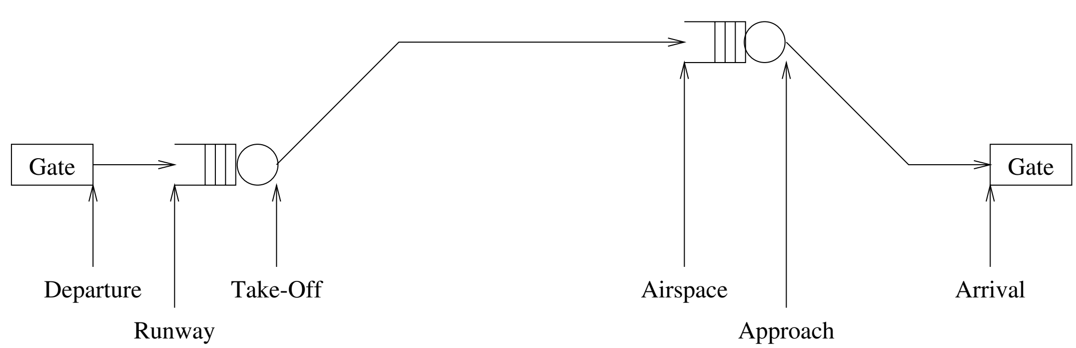


**Figure 1 Decomposition of Block Time**

After the arrival event of a leg, it is determined whether the plane requires maintenance service. The scheduled maintenance of a plane is given by the rotation. However, if a disruption occurs, the recovery policy may reschedule the maintenance inspections. Upon the arrival of a leg, the aircraft requires unscheduled maintenance service with probability $p$. The duration of the unscheduled maintenance and $p$ depend on the aircraft. A *repaired plane event* occurs when the plane is prepared to fly again.

In addition to the block time and repaired-plane events, the model includes weather and congestion events that alter the weather and congestion states at each station as described in §2.1.

### 2.3. Operational Decisions

A recovery policy consists of several *recovery components*. For example, one recovery component may define how the recovery policy will reroute disrupted crews. Each recovery component consists of a set of operational decisions. Our stochastic model may use several operational decisions in order to implement a specific recovery component. In our model,

- Legs may be delayed.
- Legs may be cancelled.
- Crews may be *deadheaded*; that is, they can fly as passengers.
- Planes may be *ferried*; that is, they can be flown to another station without passengers.
- Reserve crews may be called to fly pairings.
- Planes may be *swapped*; that is, they may be rerouted.
- Crews may be rerouted on new reconstructed pairings.
- Passengers may be rerouted.

Throughout this paper, random variables are written in **boldface** and deterministic parameters are denoted by `typewriter` font.

**2.3.1. Delays and Cancellations.** Every recovery policy must use a set of recovery components that maintain flow balance for every plane, crew, and passenger flow. Delaying a leg may not change the flow of the system, but cancelling a leg requires either rerouting the flows or additional cancellations. Consider the routing of a plane consisting of the legs in Table 3. We refer to a set of consecutive legs that begin and end at the same station as a *cancellation cycle*. If Flight 31 is cancelled, then the airlines may also cancel Flights 32 and 33, so that the plane can continue to fly out of Minneapolis. Therefore, Flights 31, 32, and 33 are a cancellation cycle. Two other cancellation cycles are Flights 33 and 34 and Flights 34 and 35. Cancelled legs are removed from the state.

**Table 3** **Plane Rotation**

| Flight | Origin | Scheduled Departure Day | Scheduled Departure Time | Destination | Scheduled Arrival Time |
|---|---|---|---|---|---|
| 31 | Minneapolis | Monday | 19:20 | San Diego | 21:40 |
| 32 | San Diego | Tuesday | 6:45 | Burbank | 7:28 |
| 33 | Burbank | Tuesday | 8:25 | Minneapolis | 14:16 |
| 34 | Minneapolis | Tuesday | 15:00 | Burbank | 17:18 |
| 35 | Burbank | Tuesday | 18:00 | Minneapolis | 23:46 |
| 36 | Minneapolis | Wednesday | 7:39 | San Juan | 14:07 |

*Note.* The cancellation cycles are (31, 32, 33), (33, 34), and (34, 35).

**2.3.2. Deadheads.** If a leg in a pairing is cancelled or flown by another crew, then the original crew can be moved from the departure station of the cancelled leg to the arrival station by deadheading. The crew is permitted to deadhead on another fleet or even another airline. We use a station-by-station matrix of legs that would be flown by another fleet or another airline carrier. Each entry of the matrix contains an origin station, a destination station, a duration, and a cost. All duration entries not supplied initially are determined using an all-pairs shortest-path algorithm with the supplied entries. At any time, a crew can start from the origin station and get to the destination station in the constant length of time specified. The crew will then incur a fixed cost, which may correspond to the price of a ticket. Airlines often have reciprocal agreements that allow crews to deadhead on other airlines for free. In this case, the fixed cost would be zero.

**2.3.3. Ferries.** Although airlines rarely ferry a plane, a station-by-station matrix similar to the one used for deadheads is used to move the planes. For example, if Flight 31 in Table 3 is cancelled, then the plane could be ferried to San Diego to fly Flight 32.

**2.3.4. Reserve Crews.** Airlines attempt to prevent a pairing from violating planning rules in operations even though the FAA allows them to do so under extreme conditions (Fedral Aviation Administration 1999). If a disruption causes a pairing to violate planning rules or if flying the planned flight schedule would cause a pairing to violate planning rules, the recovery policy could propose calling upon a reserve crew.

When the recovery policy calls upon a reserve crew, it gives a partial pairing for the reserve crew to fly. If there exists an available reserve crew, then the recovery policy finds one to assign to the legs. If the starting leg of the partial pairing departs from a station different from the crew base of the reserve crew, the recovery policy deadheads the crew from its crew base to the departure station of the starting leg. When the reserve crew finishes flying the last leg of the partial pairing, it may be necessary to deadhead the crew back to its crew base. The cost of a reserve crew is the number of minutes it flies, and the originally scheduled crew is paid the planned cost of the pairing.

**2.3.5. Plane Swaps.** A plane swap occurs when two or more planes switch their routings. The new rotations are assigned to the planes, and the legs that the planes are scheduled to fly next are determined. When cancelling legs, a controller often cancels additional legs to maintain flow balance, but plane swaps can be used to reduce the number of these additional cancellations. We present an example of how two planes may be swapped. This example and others are based upon information from an actual fleet from a major domestic carrier. Suppose Planes A and B are scheduled to fly Routes 1 and 2 in Table 4, respectively. Suppose Plane A requires 15 hours of unscheduled maintenance at 10:15 on Monday in Burbank, and so a controller must cancel Flight 12. This decision would leave Plane A in Burbank, even though its next scheduled leg is Flight 13 departing from Boise. In addition to cancelling Flight 12, the recovery policy proposes cancelling Flight 23 and rerouting Route 1 in Burbank and Route 2 in Boise on Monday afternoon. Flight 24 becomes part of Route 1, and Route 2 includes Flights 13, 14, and 15. The recovery policy schedules the planes to fly the legs as in Table 5. Consequently, Plane A continues flying Flight 24 from Burbank after its maintenance service.

**Table 4 Original Routes**

| Route | Scheduled Departure Day | Flight | Origin | Scheduled Departure Time | Destination | Scheduled Arrival Time |
|---|---|---|---|---|---|---|
| 1 | Monday | 11 | New York | 7:30 | Burbank | 10:07 |
| | Monday | 12 | Burbank | 11:25 | Boise | 14:43 |
| | Tuesday | 13 | Boise | 14:39 | San Diego | 16:00 |
| | Tuesday | 14 | San Diego | 17:40 | Boise | 20:47 |
| | Tuesday | 15 | Boise | 21:30 | Oakland | 23:03 |
| 2 | Monday | 21 | Minneapolis | 6:30 | Boston | 9:29 |
| | Monday | 22 | Boston | 11:45 | Boise | 13:59 |
| | Tuesday | 23 | Boise | 15:30 | Burbank | 16:48 |
| | Tuesday | 24 | Burbank | 18:00 | Minneapolis | 23:46 |

**2.3.6. Reconstructed Pairings.** Similar to swapping planes, pairings can be rerouted. Some crew recovery policies, such as those described in Lettovský et al. (2000) and Stojković et al. (1998), construct new pairings. Our model can include such recovery policies.

**2.3.7. Rerouting Passenger Itineraries.** When a passenger cannot make his flight connection, the model may try to reroute him on another set of legs, including the possibility of using another airline carrier.

# 3. Recovery Components

In practice, airlines implement recovery decisions manually in the AOCC. This is typically done without a well-defined recovery policy. In this section, we describe several components of different recovery

**Table 5 New Routes**

| Route | Scheduled Departure Day | Flight | Origin | Scheduled Departure Time | Destination | Scheduled Arrival Time |
|---|---|---|---|---|---|---|
| 1 | Monday | 11 | New York | 7:30 | Burbank | 10:07 |
| | Tuesday | 24 | Burbank | 18:00 | Minneapolis | 23:46 |
| 2 | Monday | 21 | Minneapolis | 6:30 | Boston | 9:29 |
| | Monday | 22 | Boston | 11:45 | Boise | 13:59 |
| | Tuesday | 13 | Boise | 14:39 | San Diego | 16:00 |
| | Tuesday | 14 | San Diego | 17:40 | Boise | 20:47 |
| | Tuesday | 15 | Boise | 21:30 | Oakland | 23:03 |

policies which use the operational decisions described in §2.3.

### 3.1. Push-Back

When a leg is delayed, the recovery policy needs to respond to the delay. The recovery policy may use a simple routine that delays a leg until its scheduled plane and crew are ready, regardless of their tardiness. We refer to this recovery component as *push-back*. The *slack* in a plane's rotation refers to the length of time from the end of a leg to the start of the next leg in the rotation. Similarly, the slack in a crew's pairing refers to the length of time between legs in the pairing. With a sufficient amount of slack in an airline system and short delays, push-back performs well. If a leg is delayed, then the assigned crew and plane have a sufficient amount of time to connect to their next legs.

**3.1.1. Compensatory Rest Delays.** Push-back also accounts for 8-in-24 planning violations that require compensatory rest for the crews. For example, consider the scheduled pairing in Table 6. The crew is briefed before each duty for 60 minutes and is debriefed after each duty for 15 minutes. In the first three days of operations, the crew flew the legs in Table 7. In the 24 hours between Tuesday at 22:56 and Wednesday at 22:56, the crew flew Flights 43, 44, and 45 for a total flying time of 9 hours and 9 minutes. The crew left for its rest in San Diego, after being debriefed after Flight 43, at 5:01. They then returned to the airport one hour before the scheduled departure of Flight 44, at 15:40. The crew had a 10-hour and 39-minute rest period. By Federal Aviation Regulation 121.471(c3), the crew must receive 12 hours of compensatory rest after Flight 45 (Federal Aviation Administration 1999). Furthermore, the crew must be debriefed after Flight 45 and briefed before Flight 46. Push-back would schedule Flight 46, guaranteeing the crew 12 hours of rest, for a departure time of 12:11 on Thursday. If Flight 46 flies as planned, the revised schedule would be as displayed in Table 8. Observe that push-back does not disrupt plane or crew flow balance.

**Table 6** **Originally Scheduled Pairing**

| Flight | Origin | Scheduled Departure Day | Scheduled Departure Time (EST) | Destination | Scheduled Arrival Time (EST) |
|---|---|---|---|---|---|
| 41 | Burbank | Monday | 18:30 | New York | 23:52 |
| 42 | New York | Tuesday | 15:45 | Minneapolis | 18:26 |
| 43 | Minneapolis | Tuesday | 20:20 | San Diego | 0:40 |
| 44 | San Diego | Wednesday | 16:40 | Burbank | 17:24 |
| 45 | Burbank | Wednesday | 19:05 | Minneapolis | 22:50 |
| 46 | Minneapolis | Thursday | 9:30 | Burbank | 13:53 |

**Table 7** **Pairing After Wednesday**

| Flight | Origin | Departure Day | Departure Time (EST) | Destination | Arrival Time (EST) |
|---|---|---|---|---|---|
| 41 | Burbank | Monday | 18:32 | New York | 23:44 |
| 42 | New York | Tuesday | 18:23 | Minneapolis | 20:59 |
| 43 | Minneapolis | Wednesday | 0:11 | San Diego | 4:46 |
| 44 | San Diego | Wednesday | 16:42 | Burbank | 17:31 |
| 45 | Burbank | Wednesday | 19:11 | Minneapolis | 22:56 |

### 3.2. Short Cycle Cancellation

The push-back recovery component may allow delays to propagate. If a plane is significantly behind schedule, the recovery policy can propose cancelling a leg. When a recovery policy cancels a leg, it must maintain flow balance for the planes, the crews, and the passengers. Because crews and passengers can fly on other legs, there are many ways to maintain a crew's and passenger's flow balance. Unfortunately, rerouting the plane is more difficult, and so the airlines may cancel additional legs on the plane's route.

Determining which cancellation cycle to cancel is not always easy. Consider the following *short cycle cancellation*. When a leg incurs a long delay, short cycle cancellation considers the cycles that start with the legs on the plane's rotation. For each cycle $c$, the recovery policy calculates the following attributes:

*Attribute 1.* The total amount of passenger revenue on cycle $c$. The passengers on cycle $c$ must be rerouted to their destinations.

*Attribute 2.* The total amount of passenger revenue from legs after the disruption and before the first leg

**Table 8** **Pairing Schedule on Thursday**

| Flight | Origin | Scheduled Departure Day | Scheduled Departure Time (EST) | Destination | Scheduled Arrival Time (EST) |
|---|---|---|---|---|---|
| 46 | Minneapolis | Thursday | 12:11 | Burbank | 16:34 |

of cycle $c$. The passengers on the legs after the disruption are inconvenienced by a long delay.

*Attribute 3.* The lateness of the plane after cycle $c$ has been cancelled, assuming no other delays. The lateness of the plane can create future delays.

Short cycle cancellation finds the cycle that minimizes a linear combination of Attributes 1, 2, and 3. For attribute $i = 1, 2, 3$, let $\alpha_i \geq 0$ be a predefined parameter. These parameters may be viewed as penalties for the attributes. Let $C$ be the set of all cancellation cycles on the plane's rotation. For each cancellation cycle $c \in C$, let $x_i(c)$ be the value of attribute $i = 1, 2, 3$. The penalty of cycle $c$ is given by

$$g(c) = \alpha_1 x_1(c) + \alpha_2 x_2(c) + \alpha_3 x_3(c). \tag{1}$$

Observe that $x_1(c)$ and $x_2(c)$ are in terms of passenger revenue, while the units of $x_3(c)$ are minutes. To calculate $g(c)$ in dollars, we let the units of $\alpha_3$ be dollars per minute. Let $c^*$ be the least-penalized cycle; that is,

$$c^* \in \arg\min_{c \in C} \{g(c)\}. \tag{2}$$

Let $\psi$ be the cost of pushing back the current schedule. If $g(c^*) < \psi$, then short cycle cancellation proposes cancelling cycle $c^*$. If $g(c^*) \geq \psi$, then push-back recovery is used rather than cancellation.

We demonstrate the short cycle cancellation with the following example. Suppose that the passenger revenue on each leg is \$10,000. Let $\alpha_1 = 10$, $\alpha_2 = 5$, and $\alpha_3 = \$0$ per minute. Table 9 displays a plane's future rotation. Before Flight 61, the plane incurs a 10-hour and 19-minute maintenance delay. The plane requires a minimum of 15 minutes between legs. Assuming the crews are always available and there are no additional delays, Table 10 displays the estimated flight times, and suppose $\psi = \$500{,}000$. Consider cancellation cycles $c_1 = (62, 63, 64, 65)$ and $c_2 = (63, 64)$. The penalty of $c_1$ is $g(c_1) = (4 \times 10 \times \$10{,}000) + (1 \times 5 \times \$10{,}000) = \$450{,}000$. The penalty of $c_2$ is $g(c_2) = (2 \times 10 \times \$10{,}000) + (2 \times 5 \times \$10{,}000) = \$300{,}000$. Because $g(c_2) < g(c_1)$ and $g(c_2) < \$500{,}000$, short cycle cancellation cancels Flights 63 and 64 instead of Flights 62, 63, 64, and 65. Table 11 displays the revised schedule after the cancellation has been made.

**Table 9** **Plane's Scheduled Rotation**

| Flight | Origin | Scheduled Departure Day | Scheduled Departure Time | Destination | Scheduled Arrival Time |
|---|---|---|---|---|---|
| 61 | Oakland | Monday | 16:15 | New York | 0:46 |
| 62 | New York | Tuesday | 6:45 | Pittsburgh | 8:03 |
| 63 | Pittsburgh | Tuesday | 8:55 | Portland | 12:28 |
| 64 | Portland | Tuesday | 13:00 | Pittsburgh | 20:41 |
| 65 | Pittsburgh | Tuesday | 21:35 | New York | 22:55 |
| 66 | New York | Wednesday | 9:35 | Minneapolis | 11:13 |
| 67 | Minneapolis | Wednesday | 12:00 | Burbank | 14:19 |

**Table 10** **Estimation of Flight Times**

| Flight | Origin | Estimated Departure Day | Estimated Departure Time | Destination | Estimated Arrival Time |
|---|---|---|---|---|---|
| 61 | Oakland | Tuesday | 2:34 | New York | 11:05 |
| 62 | New York | Tuesday | 11:20 | Pittsburgh | 12:38 |
| 63 | Pittsburgh | Tuesday | 12:53 | Portland | 16:26 |
| 64 | Portland | Tuesday | 16:41 | Pittsburgh | 0:22 |
| 65 | Pittsburgh | Tuesday | 0:37 | New York | 1:57 |
| 66 | New York | Wednesday | 9:35 | Minneapolis | 11:13 |
| 67 | Minneapolis | Wednesday | 12:00 | Burbank | 14:19 |

*Note.* Flight 61 incurs a 10-hour and 19-minute unscheduled maintenance delay. We assume there are no additional delays, the crews are always available, and the plane requires 15 minutes between legs.

### 3.3. Reserve Crews for Planning Violations

Whenever a crew is violating planning rules or is scheduled to violate them, a reserve crew can be called. For example, consider the scheduled plan for the pairing in Table 12. Due to an unscheduled aircraft maintenance delay, Flight 71 departed 3 hours late, and so Flight 72 departed at 17:16, 2 hours and 36 minutes late, and Flight 73 departed at 19:33, 2 hours and 18 minutes late. The crew was debriefed after Flight 73 for 15 minutes and was briefed before Flight 74 for one hour. Table 13 displays the times

**Table 11** **Results of the Short Cycle Cancellation**

| Flight | Origin | Scheduled Departure Day | Scheduled Departure Time | Destination | Scheduled Arrival Time |
|---|---|---|---|---|---|
| 61 | Oakland | Tuesday | 2:34 | New York | 11:05 |
| 62 | New York | Tuesday | 11:20 | Pittsburgh | 12:38 |
| 65 | Pittsburgh | Tuesday | 21:35 | New York | 22:55 |
| 66 | New York | Wednesday | 9:35 | Minneapolis | 11:13 |
| 67 | Minneapolis | Wednesday | 12:00 | Burbank | 14:19 |

*Note.* Cycle (63, 64) is cancelled.

**Table 12 Scheduled Pairing**

| Flight | Origin | Scheduled Departure Day | Scheduled Departure Time (EST) | Destination | Scheduled Arrival Time (EST) |
|---|---|---|---|---|---|
| 71 | Minneapolis | Monday | 12:05 | Birmingham | 13:49 |
| 72 | Birmingham | Monday | 14:40 | Washington, D.C. | 16:15 |
| 73 | Washington, D.C. | Monday | 17:15 | New York | 18:48 |
| 74 | New York | Tuesday | 6:30 | Washington, D.C. | 8:03 |
| 75 | Washington, D.C. | Tuesday | 8:55 | Portland | 14:28 |
| 76 | Portland | Wednesday | 13:00 | Burbank | 15:39 |
| 77 | Burbank | Wednesday | 20:00 | Minneapolis | 23:45 |

for the legs after Flight 74 arrived at the gate. In the 24-hour period between Monday at 8:10 and Tuesday at 8:10, the crew flew Flights 71, 72, 73, and 74 for a total flying time of 6 hours and 25 minutes. The crew did not leave the airport after Flight 73 until 21:11 (15 minutes after 20:56), and the crew arrived at the airport on Tuesday at 5:30 (one hour before 6:30). Thus, the crew had 8 hours and 19 minutes of rest, and so the pairing does not violate 8-in-24 planning rules and does not need compensatory rest.

Consider the pairing as if it flies Flight 75 as scheduled. In the 24-hour period between Monday at 14:28 and Tuesday at 14:28, the crew would have flown flights 71, 72, 73, 74, and 75 for a total flying time of 11 hours and 58 minutes. By Federal Aviation Regulation Section 121.471 (c3) (Federal Aviation Administration 1999), this crew would violate 8-in-24 planning rules. If there is a reserve crew available, the recovery policy calls a reserve crew to fly Flights 75, 76, and 77. However, if there are no reserve crews available, the recovery policy determines another recovery solution.

**3.3.1. First Available Reserve Crew Selection.** When the recovery policy proposes that a reserve crew fly a partial pairing, it must also select a specific reserve crew. We propose the following *first available reserve crew selection* that minimizes a weighted sum of:

*Attribute 1.* The length of time required to wait for the reserve crew to becomes available;

*Attribute 2.* The length of time required to deadhead the reserve crew from the crew base to the origin station of the first leg of the partial pairing if necessary;

*Attribute 3.* The length of time required to deadhead the reserve crew from the destination station of the last leg of the partial pairing to the crew base if necessary.

Let $\beta_1$, $\beta_2$, and $\beta_3$ be the penalties of Attributes 1, 2, and 3, respectively, and for each reserve crew, the penalized sum of these attributes is called the *excess time*.

Let **currenttime** be the current time of the stochastic process. For any two stations $s_1$ and $s_2$, let $\texttt{deadheadtime}(s_1, s_2)$ be the length of time required to deadhead a crew from station $s_1$ to station $s_2$. For each flight $f$, let $\texttt{departurestation}(f)$ and $\texttt{arrivalstation}(f)$ be the origin and destination station of leg $f$. Let $\{f_1, \ldots, f_n\}$ be the sequence of legs in the partial pairing. For each reserve crew $\lambda$, let

**Table 13 Pairing After Flight 74 Arrived at the Gate**

| Flight | Origin | Departure Day | Departure Time (EST) | Destination | Arrival Time (EST) |
|---|---|---|---|---|---|
| 71 | Minneapolis | Monday | 15:05 | Birmingham | 16:59 |
| 72 | Birmingham | Monday | 17:16 | Washington, D.C. | 18:46 |
| 73 | Washington, D.C. | Monday | 19:33 | New York | 20:56 |
| 74 | New York | Tuesday | 6:32 | Washington, D.C. | 8:10 |

**reservecrewtime**($\lambda$) be the time reserve crew $\lambda$ will be available, let `crewbase`($\lambda$) be the crew base for reserve crew $\lambda$, and let **excess**($\lambda$) be the excess time for reserve crew $\lambda$, which is given by

$$\begin{aligned}\textbf{excess}(\lambda) = \beta_1 \times \max\{&\textbf{reservecrewtime}(\lambda) \\ &-\textbf{currenttime}, 0\} \\ +\beta_2 \times &\,\texttt{deadheadtime}(\texttt{crewbase}(\lambda), \\ &\texttt{departurestation}(f_1)) \\ +\beta_3 \times &\,\texttt{deadheadtime}(\texttt{arrivalstation}(f_n), \\ &\texttt{crewbase}(\lambda)).\end{aligned} \tag{3}$$

Because the recovery policy prefers to limit the number of times it deadheads crews, the reserve crew selection tries to find the reserve crew with the smallest amount of **excess** at either the departure station of the first leg or the arrival station of the last leg of the partial pairing. Let $\Lambda(s)$ be the set of available crews at station $s$. If $s$ is not a crew base, then $\Lambda(s) = \emptyset$. Let $\varepsilon$ be a nonnegative time tolerance defined by the recovery policy that limits the excess time of the reserve crew. Let $\Lambda_1$ be the set of available reserve crews that have a minimum excess time and are stationed at either the departure station of the first leg or the arrival station of the last leg; that is,

$$\begin{aligned}\Lambda_1 = \arg\min_{\lambda}\{&\textbf{excess}(\lambda) \mid \textbf{excess}(\lambda) < \varepsilon, \\ &\lambda \in \Lambda(\texttt{departurestation}(f_1)) \\ &\cup \Lambda(\texttt{arrivalstation}(f_n))\}.\end{aligned} \tag{4}$$

If $\Lambda_1 \neq \emptyset$, the recovery policy calls a reserve crew $\lambda^* \in \Lambda_1$. If there are no available crews at the departure station of the first leg or the arrival station of the last leg, $\Lambda_1 = \emptyset$, then the reserve crew selection searches all of the crew bases for an available crew. Let $B$ be the set of all crew bases, and let $\Lambda_2$ be the set of all available reserve crews that have a minimum excess time:

$$\Lambda_2 = \arg\min_{\lambda}\{\textbf{excess}(\lambda) \mid \textbf{excess}(\lambda) < \varepsilon, \lambda \in \Lambda(s), s \in B\}. \tag{5}$$

If $\Lambda_2 \neq \emptyset$, then it calls a reserve crew, $\lambda^* \in \Lambda_2$. Otherwise, the first available reserve crew selection determines that there is no available reserve crew with an excess time less than $\varepsilon$, and so the recovery policy tries another method to overcome the disruption.

### 3.4. Passenger Push-Back

Airlines will sometimes delay legs so that passengers will not miss their connecting legs. For example, if in a passenger itinerary leg $f_1$ connects to leg $f_2$, and leg $f_1$ is delayed, then a controller would delay leg $f_2$ so that the passengers could make their connection. Consider the following *passenger push-back* recovery component. Let `minmisconnections` be a nonnegative integer. Passenger push-back proposes a delay if at least `minmisconnections` passengers would miss their connections, assuming there are no other delays or cancellations. Let `passengerdelay` be the number of minutes passenger push-back proposes delaying the leg. Let EPM($f, d$) be the estimated number of passengers that would miss their connections if leg $f$ is delayed by $d$ minutes and there are no subsequent delays or cancellations. Therefore, if

$$\begin{aligned}\text{EPM}(f, 0) > \text{EPM}(&f, \texttt{passengerdelay}) \\ &+\texttt{minmisconnections},\end{aligned} \tag{6}$$

then passenger push-back would propose delaying leg $f$ by `passengerdelay` minutes.

We demonstrate passenger push-back using the following example. Let `passengerdelay` = 10 minutes and `minmisconnections` = 50 passengers, and assume that passengers need at least 20 minutes to make their connections if they are changing planes. Table 14 displays the itineraries that include Flight 85. Suppose Flight 81 departed as planned and arrived three minutes early. Due to a snow storm in Boise, Flight 82 departed 26 minutes late and is scheduled to land 53 minutes late. Flight 83 incurred a seven-hour unscheduled maintenance delay. Flight 84 departed on time but arrived 15 minutes late. Table 15 displays the itinerary schedule at 18:26.

Observe that the passengers on Itineraries 2 and 5 will make their connecting legs, but passengers on Itineraries 3 and 4 will miss their connections, and so EPM($f$, 0) = 68. The passengers on Itinerary 3 only need 10 additional minutes between Flights 82 and 85 to have the 20 minutes needed to make their connection, so EPM($f$, 10) = 10. Passenger push-back would suggest delaying Flight 85 by 10 minutes so that the 58 passengers on Itinerary 3 could make their connection. Table 16 shows the resulting itinerary schedule.

**Table 14 Scheduled Itineraries**

| Itinerary | Number of Passengers | Flight | Origin | Scheduled Departure Time | Destination | Scheduled Arrival Time | Plane |
|---|---|---|---|---|---|---|---|
| 1 | 34 | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |
| 2 | 45 | 81 | New York | 16:45 | Minneapolis | 18:29 | A |
| | | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |
| 3 | 58 | 82 | Boise | 15:00 | Minneapolis | 18:17 | B |
| | | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |
| 4 | 10 | 83 | Newark | 16:30 | Minneapolis | 17:43 | C |
| | | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |
| 5 | 3 | 84 | New York | 15:45 | Minneapolis | 17:26 | D |
| | | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |

*Note.* Passengers need at least 20 minutes to make their connections if they are changing planes.

**Table 15 Scheduled Itineraries at 18:26**

| Itinerary | Number of Passengers | Flight | Origin | Departure Time | Destination | Arrival Time | Plane |
|---|---|---|---|---|---|---|---|
| 1 | 34 | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |
| 2 | 45 | 81 | New York | 16:45 | Minneapolis | 18:26 | A |
| | | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |
| 3 | 58 | 82 | Boise | 15:26 | Minneapolis | 19:10 | B |
| | | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |
| 4 | 10 | 83 | Newark | 23:30 | Minneapolis | 0:43 | C |
| | | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |
| 5 | 3 | 84 | New York | 15:45 | Minneapolis | 17:41 | D |
| | | 85 | Minneapolis | 19:20 | Burbank | 21:40 | A |

*Note.* Passengers on Itineraries 3 and 4 will miss their connections.

**Table 16 Resulting Itinerary Schedule After Passenger Push-Back**

| Itinerary | Number of Passengers | Flight | Origin | Departure Time | Destination | Arrival Time | Plane |
|---|---|---|---|---|---|---|---|
| 1 | 34 | 85 | Minneapolis | 19:30 | Burbank | 21:50 | A |
| 2 | 45 | 81 | New York | 16:45 | Minneapolis | 18:26 | A |
| | | 85 | Minneapolis | 19:30 | Burbank | 21:50 | A |
| 3 | 58 | 82 | Boise | 15:26 | Minneapolis | 19:10 | B |
| | | 85 | Minneapolis | 19:30 | Burbank | 21:50 | A |
| 4 | 10 | 83 | Newark | 23:30 | Minneapolis | 0:43 | C |
| | | 85 | Minneapolis | 19:30 | Burbank | 21:50 | A |
| 5 | 3 | 84 | New York | 15:45 | Minneapolis | 17:41 | D |
| | | 85 | Minneapolis | 19:30 | Burbank | 21:50 | A |

*Note.* By delaying Flight 85 by 10 minutes, the passengers on Itinerary 3 make their connection.

Observe that only the 10 passengers on Itinerary 4 will miss their connection.

## 4. Performance Measures

There are many criteria that can be used to evaluate the quality of a schedule. Most airlines pay crews according to the maximum of their planned pairing cost and their operational pairing cost. In addition, reserve crews are expensive. Results in operations, such as on-time performance and cancellations, affect the airline's image. Misconnecting and cancelled itineraries reduce customer satisfaction.

### 4.1. Crew Cost and FTC

Crews are paid proportionally to the number of *pay-and-credit* minutes they accumulate. We refer to the number of pay-and-credit minutes that an individual crew accumulates as its *crew cost*. Before we describe the actual cost of a pairing, we explain its planned cost by formulas that are used by major domestic airlines.

For each leg $f$, let $\texttt{originalblock}(f)$ be the originally planned block time of leg $f$. Let $\texttt{elapse}_d$ be the planned elapsed time of duty $d$. Let $r_e$ be a fraction representing the rate of pay for the elapsed time, and let *mgd* be the minimum guarantee for a duty. The *planned duty cost* of duty $d$ is assumed to be

$$\texttt{b}_d = \max\left\{\sum_{f\in d}\texttt{originalblock}(f),\; r_e \times \texttt{elapse}_d,\; mgd\right\}. \tag{7}$$

Let $\texttt{TAFB}_p$ be the planned time away from base of pairing $p$. Let $r_t$ be a fraction representing the rate of pay of TAFB. Let *mgp* be a minimum guarantee per duty in a pairing, and let $numduties_p$ be the number of duties in pairing $p$. Then the *planned pairing cost* of pairing $p$ is

$$\texttt{c}_p = \max\left\{\sum_{d\in p}\texttt{b}_d,\; r_t \times \texttt{TAFB}_p,\; mgp \times numduties_p\right\}. \tag{8}$$

Vance et al. (1997) use values of $r_e = 4/7$, $mgd = 0$, $r_t = 2/7$, and $mgp = 300$ in their branch-and-price heuristic for crew scheduling.

The actual cost, as well as the planned cost, of a pairing depends on the contractual agreement between an airline and the pilots. We describe one method for calculating the *actual pairing cost* used by a major domestic carrier. The operational cost of a pairing is calculated in the same way as the planned cost of the pairing, but using the actual departure and arrival times instead of the scheduled times. That is, the *operational duty cost* of a duty $d$, $\mathbf{b}_d$, is given by

$$\mathbf{b}_d = \max\left\{\sum_{f\in d}\mathbf{block}(f),\; r_e \times \mathbf{elapse}_d,\; mgd\right\}, \tag{9}$$

where $\mathbf{block}(f)$ is the *operational block time* of leg $f$, and $\mathbf{elapse}_d$ is the *operational elapsed time* of duty $d$. The *operational pairing cost* of pairing $p$ is given by

$$\mathbf{c}_p = \max\left\{\sum_{d\in p}\mathbf{b}_d,\; r_t \times \mathbf{TAFB}_p,\; mgp \times numduties_p\right\}, \tag{10}$$

where $\mathbf{TAFB}_p$ is the *operational time away from base* of pairing $p$. The *actual pairing cost* of the pairing $p$ is

$$\tilde{\mathbf{c}}_p = \max\{\texttt{c}_p, \mathbf{c}_p\}. \tag{11}$$

Let $P$ be the set of all pairings in a crew schedule, and so the actual number of pay-and-credit minutes is

$$\sum_{p\in P}\tilde{\mathbf{c}}_p.$$

Airlines often use flight-time credit (FTC) as a measurement of the cost of a crew schedule. FTC is the difference between the number of minutes paid and the number of minutes flown, as a percentage of the number of minutes flown (Vance et al. 1997):

$$\text{FTC} = \frac{\text{pay-and-credit minutes} - \text{flytime}}{\text{flytime}} \times 100\%. \tag{12}$$

### 4.2. Cancellations and On-Time Percentage

According to the Bureau of Transportation Statistics, a leg is considered *on time* if it arrives at the gate within 15 minutes of its originally scheduled arrival time. The *on-time percentage* is the number of on-time legs as a percentage of the number of legs scheduled (Bureau of Transportation Statistics 1998). The

Bureau of Transportation Statistics counts a cancelled leg against the on-time percentage. The percentage of legs within 15 minutes and 60 minutes of their scheduled arrival time is tabulated, and the number of cancelled legs per day is recorded.

### 4.3. Passenger Misconnections

A *passenger misconnection* is a passenger who cannot fly her planned itinerary either because the itinerary has missed a connection or has been cancelled.

The number of passengers whose itinerary has either missed a connection or been cancelled is counted. For each leg $f$, let **departuretime**$(f)$ be the departure time, let **arrivaltime**$(f)$ be the arrival time, and let **nextplaneleg**$(f)$ be the leg following leg $f$ on the assigned plane's rotation. Let `minpassturn` be the time it takes passengers to change planes. Let $\kappa$ be the set of cancelled legs. Let $I$ be the set of all passenger itineraries. For each itinerary $i \in I$, let $\{f_1(i), \ldots, f_{n(i)}(i)\}$ be the sequence of legs in $i$, and let **passengers**$(i)$ be the number of passengers flying itinerary $i$. The set of *cancelled itineraries*, $K$, is the set of the itineraries in which at least one leg is cancelled,

$$K = \{i \in I \mid \exists j \in 1, \ldots, n(i), f_j(i) \in \kappa\}. \tag{13}$$

The set of *misconnecting itineraries*, $M$, consists of the uncancelled itineraries in which the passengers do not have sufficient time to make one or more of their connections when changing planes; that is,

$$\begin{aligned} M = \{i \in I - K \mid \exists j \in 1, \ldots, n(i) - 1, \\ \mathbf{nextplaneleg}(f_j(i)) \neq f_{j+1}(i), \\ \mathbf{arrivaltime}(f_j(i)) + \mathtt{minpassturn} \\ > \mathbf{departuretime}(f_{j+1}(i))\}. \end{aligned} \tag{14}$$

Observe that if the passengers stay on the same plane, $\mathbf{nextplaneleg}(f_j(i)) = f_{j+1}(i)$, then they will not miss this connection, and so we ignore them when determining $M$. The number of passenger misconnections, PM, is

$$PM = \sum_{i \in M \cup K} \mathbf{passengers}(i). \tag{15}$$

This method does not consider passengers who originally misconnected but were successfully rerouted to their destination.

### 4.4. Crew Legality and Reserve Crews

When a crew violates planning rules, the recovery policy calls upon a reserve crew to fly the remainder of the pairing and reroutes the crew. Reserve crew usage and the number of crews that violate planning rules are measures of the quality of a crew schedule and recovery policy. As a result, the number of reserve crews used per day and the number of crews that violate planning rules per day is tabulated

## 5. SimAir

The primary purpose of our stochastic model is to evaluate plans and recovery procedures. We present SimAir, a simulation implementation of the model. Because different airlines use different recovery procedures, SimAir has a flexible modular structure which is capable of handling several recovery policies. SimAir also provides recovery tools to allow external recovery algorithms to be seamlessly integrated.

SimAir contains two modules for decision making. The *Controller Module* maintains the state of the simulation. We have designed the Controller to emulate the AOCC in the sense that it recognizes disruptions and implements recovery policies. If a disruption will prevent the legs from flying as planned, the Controller requests a proposed reaction from the *Recovery Module*. SimAir's current Recovery Module uses those recoveries described in §3, even though the user can alter it to support other recovery procedures.

SimAir uses an *Event Generator Module*, which samples random ground time delays, additional block time delays, and unscheduled maintenance delays as described in §2.2. Because different airlines experience different delays, the user can easily update the Event Generator for alternate delay distributions.

Figure 2 gives a schematic representation of the structure of SimAir.

SimAir uses a *simulation clock* and a time-sorted *event queue*. There are three types of events—arrivals, departures, and repaired planes. In the implementation described in this paper, we do not consider dependent delays and restricted flow. The simulation clock is the time currently being simulated. SimAir keeps track of the first event, the last event, and the most recently added event in the event queue. These events drive the simulation. SimAir removes the first

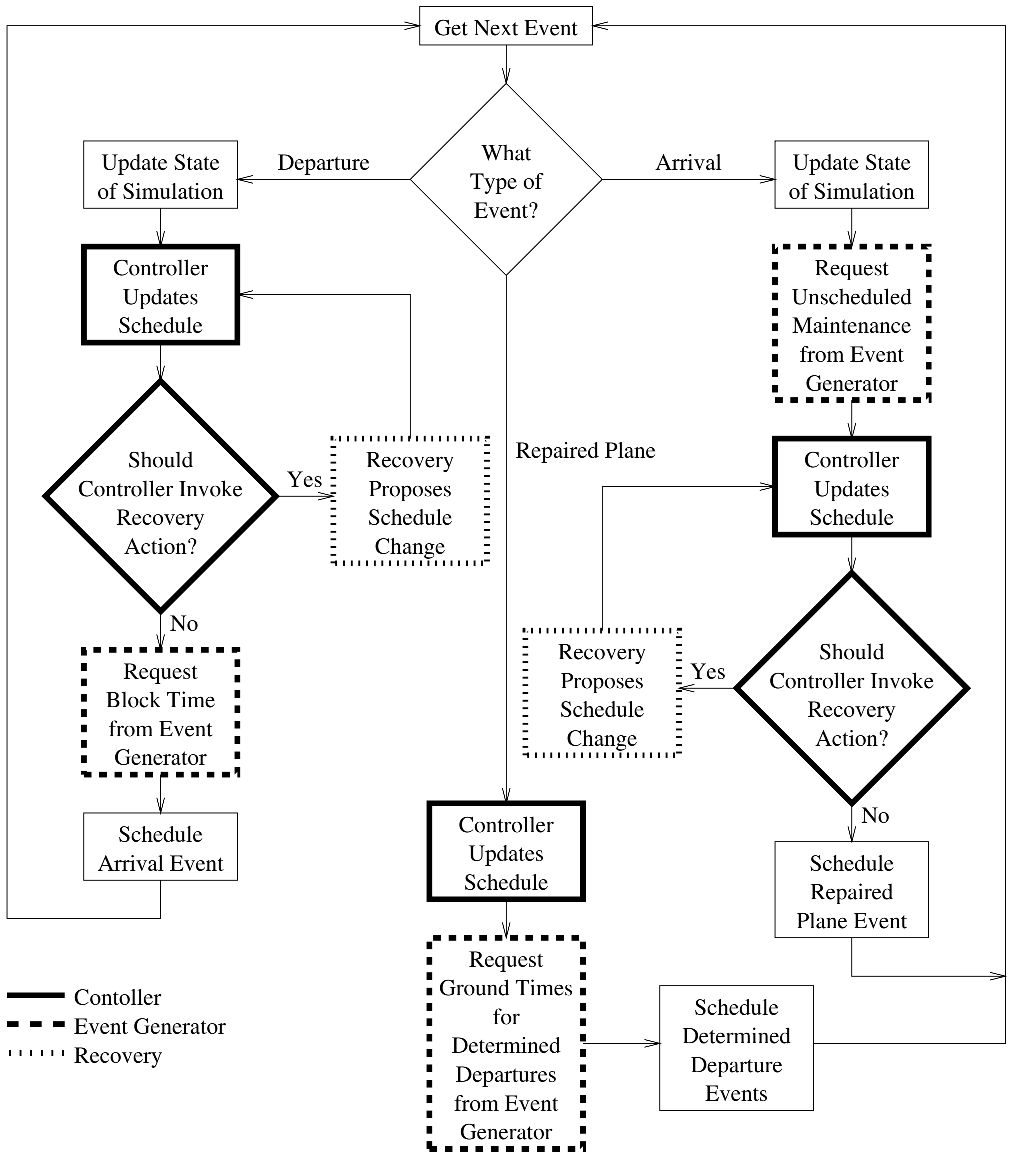


**Figure 2** **The Structure of SimAir**

event from the event queue and updates the simulation clock. SimAir can insert an event into the event queue. For example, if the first event is a departure event, then SimAir would update its simulation clock to the departure time and add an arrival event for the corresponding leg to the event queue. SimAir may also delete events from the event queue. The purpose for deleting events is recovery.

## 6. Example Test Instances

We simulated 10,000 days of operations for a fleet consisting of 119 daily legs from a major domestic airline carrier. The primary purpose of our examples was to test four different crew schedules with several recovery heuristics. The recovery policies did not support catastrophic disruptions such as severe weather conditions and major unscheduled maintenance problems. Consequently, we did not simulate such events, even though SimAir is capable of doing so.

### 6.1. Probability Distributions

For the probability distributions, the Event Generator used empirical distributions from operational data for the simulated fleet over a six-month period.

**6.1.1. Block Distribution.** For every leg $f$, the random variable $\mathbf{block}(f)$ was composed of two parts, the originally scheduled block time and a block time error. The Event Generator used three different distributions, $\mathbf{block}_1$, $\mathbf{block}_2$, and $\mathbf{block}_3$, for the block time error. It used the originally scheduled block time to determine which distribution should generate the block time error. Each distribution was independent of `originalblock`. Table 17 shows how the Event Generator obtained $\mathbf{block}(f)$ from $\texttt{originalblock}(f)$.

**6.1.2. Ground Distribution.** The ground time empirical data for the simulated fleet does not include the location and time of day. We assumed that **ground** was independent of the location and time of day of the departure event for the test instances, although this would be easy to change if the data were available.

**6.1.3. Unscheduled Maintenance.** For the test instances, we assumed that both unscheduled maintenance random variables discussed in §2.2 were independent of aircraft. We estimated the probability of a maintenance delay by

$$p = [\text{number of maintenance problems observed in time period}] \Big/ [\text{number of legs flown in time period}]. \qquad (16)$$

The range of the data for the length of service time is very large. In the empirical data, the unscheduled maintenance time ranges from one minute to 10 days. There are no instances for many values in this range, so the Event Generator linearly interpolated between known data points. We used six months of historical data to obtain the unscheduled maintenance distribution. Because we were not simulating severe maintenance problems, we limited the length of the delays to six hours. The unscheduled maintenance distribution was the same for all planes.

**Table 17** **Determining Block Time Distribution.**

| $\texttt{originalblock}(f)$(minutes) | $\mathbf{block}(f)$ |
|---|---|
| 0–119 | $\texttt{originalblock}(f) + \mathbf{block}_1$ |
| 120–239 | $\texttt{originalblock}(f) + \mathbf{block}_2$ |
| 240+ | $\texttt{originalblock}(f) + \mathbf{block}_3$ |

### 6.2. Crew Schedules

There typically are millions of legal pairings that can be used to build a crew schedule, and so there may be billions of possible crew schedules. We considered four crew schedules that were each optimal for a different objective function.

- Crew Schedule S1 was the optimal crew schedule found by minimizing the planned pairing costs; i.e. we assumed that all legs would fly as planned. This is the traditional method of selecting a crew schedule.
- Crew Schedule S2 was the optimal crew schedule found by augmenting the planned pairing costs by a penalty. For any pairing that included a duty with an elapsed time greater than a certain threshold, the crew cost was increased by an amount proportional to the difference between the elapsed time and the threshold.
- Crew Schedule S3 was the optimal crew schedule found by estimating the expected cost of a pairing using Monte Carlo sampling from delay distributions and then minimizing the expected cost. For each pairing we assumed that the planes were always available, and we simulated the cost of the pairing in operations.
- Crew Schedule S4 was the optimal crew schedule found when the pairing costs were calculated in the following manner. We assumed that only push-back recovery would be used, and that the planes would always be available. We discretized time into 15-minute intervals, and for each duty we calculated a probability distribution for the flying time and the elapsed time. The expected cost of each duty was calculated, and from this information the expected cost of each pairing was estimated.

For details on the objective functions used to construct these crew schedules, see Schaefer (2000) and Schaefer et al. (2001).

### 6.3. Tested Recovery Policies

In our computational study, we examined the push-back recovery heuristic. The Controller called a reserve crew whenever a crew violated or was scheduled to violate 8-in-24 planning rules on the current or next scheduled flight. The computational study also included three additional recovery components. In the first recovery component, the Controller guaranteed each crew nine hours of rest every night instead of the usual eight hours. In other computational tests, the Controller guaranteed the crews only eight hours of rest each night. Because the crews received sufficient rest each night, they never violated 8-in-24 planning rules. The second recovery component was passenger push-back. We studied the cases in which a leg was delayed 15 minutes and 30 minutes to wait for passengers. The third recovery component used short cycle cancellation. SimAir provided results for all 12 combinations of these additional recovery components for each of the four aforementioned crew schedules. In Table 18, we list five computational examples: A, B, C, D, and E described in §6.4.

### 6.4. Example Results

In Tables 19 through 23, we abbreviate the crew schedule as "Crew Sched," the deterministic FTC of the schedule as "Det FTC," and the average daily simulated FTC as "FTC $\mu$." We denote the variance of the simulated FTC as "FTC $\sigma^2$," the percentage of arrivals within 15 minutes of the scheduled arrival time as "OT+15," the percentage of arrivals within 60 minutes of the scheduled arrival time as "OT+60," the average number of crews violating 8-in-24 planning rules per day as "8-in-24 Vio," the average number of reserve crew calls per day as "Res Crews," the percentage of passenger misconnections as "Miss Pass %," and the average number of cancellations per day as "Can" whenever the recovery policy makes cancellations.

**Table 18** **Description of Recovery Policies Used in Computational Examples A, B, C, D, and E**

| Example | Table of Results | Minimum Overnight (Hours) | Passenger Push-back (Minutes) | Short Cycle Cancellation |
|---|---|---|---|---|
| A | 19 | 8 | 0 | No |
| B | 20 | 9 | 0 | No |
| C | 21 | 9 | 15 | No |
| D | 22 | 9 | 30 | No |
| E | 23 | 9 | 30 | Yes |

**Table 19** **The Computational Results for Example A**

| Crew Sched | Det FTC | FTC $\mu$ | FTC $\sigma^2$ | OT +15 | OT +60 | 8-in-24 Vio | Res Crews | Miss Pass % |
|---|---|---|---|---|---|---|---|---|
| S1 | 2.98 | 4.57 | 1.54 | 78.1 | 96.3 | 0.0108 | 0.25 | 0.7146 |
| S2 | 3.19 | 4.72 | 1.51 | 78.3 | 96.3 | 0.0082 | 0.23 | 0.7159 |
| S3 | 3.10 | 4.46 | 1.50 | 78.3 | 96.4 | 0.0097 | 0.22 | 0.7081 |
| S4 | 3.18 | 4.77 | 2.85 | 77.8 | 96.1 | 0.0151 | 0.27 | 0.6641 |

*Note.* The recovery policy was push-back with reserve crews when a pairing violated or was expected to violate 8-in-24 planning rules. The Controller granted the crews eight hours of rest each night.

Examples A and B demonstrate the importance of crew rest. Table 19 displays results of Example A, which used push-back and assumed the crews received at least eight hours of rest each night. Example B, shown in Table 20, considered the case in which the Controller guaranteed nine hours of rest for each crew. There were no reserve crew calls in Example B because a crew cannot violate 8-in-24 planning rules with at least nine hours of rest. The results from Example B are significantly better in crew cost than those of Example A, largely due to the fact that there was no reserve crew usage. We also tested a third recovery component in which the Controller granted the crews nine hours of rest only if the crew was scheduled to violate 8-in-24 planning rules the next day. This recovery approach performed only slightly better than Example A, but it performed worse than Example B.

The results of Examples B, C, and D, displayed in Tables 20, 21, and 22, demonstrate the effect of passenger push-back. The Controller guaranteed nine

**Table 20** **The Computational Results for Example B**

| Crew Sched | Det FTC | FTC $\mu$ | FTC $\sigma^2$ | OT +15 | OT +60 | Miss Pass % |
|---|---|---|---|---|---|---|
| S1 | 2.98 | 4.09 | 0.16 | 77.9 | 96.2 | 0.7467 |
| S2 | 3.19 | 4.24 | 0.14 | 78.1 | 96.2 | 0.7409 |
| S3 | 3.10 | 4.00 | 0.12 | 78.1 | 96.3 | 0.7354 |
| S4 | 3.18 | 4.03 | 0.12 | 77.4 | 95.9 | 0.7002 |

*Note.* The recovery policy was push-back. The Controller guaranteed nine hours rest to the crews. By doing so, the crews could not violate 8-in-24 planning rules.

**Table 21 The Computational Results for Example C**

| Crew Sched | Det FTC | FTC $\mu$ | FTC $\sigma^2$ | OT +15 | OT +60 | Miss Pass % |
|---|---|---|---|---|---|---|
| S1 | 2.98 | 4.09 | 0.16 | 77.7 | 96.2 | 0.6442 |
| S2 | 3.19 | 4.24 | 0.14 | 77.9 | 96.2 | 0.6214 |
| S3 | 3.10 | 4.00 | 0.12 | 77.9 | 96.2 | 0.6137 |
| S4 | 3.18 | 4.03 | 0.12 | 77.3 | 95.9 | 0.5841 |

*Note.* The recovery policy used push-back, and delayed legs by 15 minutes for passengers. The Controller guaranteed nine hours rest to the crews, so the crews could not violate 8-in-24 planning rules.

hours of rest to each crew, ensuring that no crew would violate 8-in-24 planning rules. The Controller also delayed legs by 15 and 30 minutes for passengers in Examples C and D, respectively. As the delays for passengers increased, the number of passenger misconnections decreased, and crew cost experienced insignificant changes. The on-time performance percentages decreased slightly as the Controller delayed the legs longer. We noticed the same effect on every example that increased passenger delays to 15 and 30 minutes. We experimented with increasing the passenger-induced delays beyond 30 minutes. The number of passenger misconnections continued to decrease, but the decrease was smaller. With delays longer than 30 minutes, the crew cost began to increase significantly.

Table 23 displays the effects of short cycle cancellation in Example E. We set $\alpha_1 = 10$, $\alpha_2 = 3$, and $\alpha_3 =$ \$0 per minute. The Controller also guaranteed nine hours of rest to the crews and delayed legs by 30 minutes to wait for passengers. From Example D to Example E, observe that the number of passenger misconnections increased because some of the legs were cancelled. There was minimal change in crew cost and

**Table 22 The Computational Results for Example D**

| Crew Sched | Det FTC | FTC $\mu$ | FTC $\sigma^2$ | OT +15 | OT +60 | Miss Pass % |
|---|---|---|---|---|---|---|
| S1 | 2.98 | 4.10 | 0.16 | 77.3 | 96.1 | 0.5736 |
| S2 | 3.19 | 4.25 | 0.14 | 77.5 | 96.2 | 0.5399 |
| S3 | 3.10 | 4.00 | 0.12 | 77.4 | 96.2 | 0.5324 |
| S4 | 3.18 | 4.03 | 0.12 | 76.8 | 95.8 | 0.4977 |

*Note.* The recovery policy used push-back, and delayed legs by 30 minutes for passengers. The Controller guaranteed nine hours rest to the crews, so the crews could not violate 8-in-24 planning rules.

**Table 23 The Computational Results for Example E**

| Crew Sched | Det FTC | FTC $\mu$ | FTC $\sigma^2$ | OT +15 | OT +60 | Can | Miss Pass % |
|---|---|---|---|---|---|---|---|
| S1 | 2.98 | 4.12 | 0.27 | 77.3 | 96.1 | 0.10 | 0.6722 |
| S2 | 3.19 | 4.26 | 0.21 | 77.5 | 96.2 | 0.11 | 0.6409 |
| S3 | 3.10 | 4.02 | 0.41 | 77.4 | 96.2 | 0.09 | 0.6219 |
| S4 | 3.18 | 4.05 | 0.35 | 76.8 | 95.8 | 0.10 | 0.5931 |

*Note.* The recovery policy used push-back with short cycle cancellation, and delayed legs by 30 minutes for passengers. The Controller guaranteed nine hours rest to the crews, so the crews could not violate 8-in-24 planning rules.

on-time percentage. We found this to be the case for all test instances that used short cycle cancellation. Because the tests did not include severe disruptions, there were very few cases in which the planes were significantly behind schedule, and so very few legs were cancelled. We believe that short cycle cancellation would perform better if more severe disruptions were simulated.

Finally, observe that Crew Schedule S3 outperformed S1 in each of the examples presented, even though S1 was the schedule found using state-of-the-art deterministic methods. This suggests that crew-scheduling algorithms that consider uncertainty may perform better in operations.

## 7. Conclusions and Future Research

One limitation of SimAir is that we only consider independent events. As mentioned in the description of the model, weather and airport congestion can be dependent among stations. For example, a storm in Chicago may move to Detroit. Since this paper has been written, we have implemented weather and airport congestion disruptions.

There have been very few studies on automated recovery policies. Most recovery decisions are made manually by an AOCC. Because recovery decisions have a significant effect on operations and profit, a sophisticated automated recovery module should be beneficial. Our stochastic model provides a suitable environment for the study of recovery policies in operations. Moreover, the model can assist in developing airline planning models. Many planning models are solved using optimization models. Most of these models assume every leg flies as planned. Because airline operations rarely follow the initial plan, considering disruptions may lead to plans that perform better in practice. The model provides a more realistic environment to measure the performance of an airline plan in operations.

In this paper, we provide examples using simple recovery heuristics. We could easily adjust these heuristics to test additional policies, such as relocating reserve crews. Although the recovery components described in this paper do not include planes changing rotations, the airlines often make such adjustments, and we recently implemented an aircraft recovery that swaps planes as described in Rosenberger et al. (2001). In addition to plane changes, recent optimization techniques that reroute crew pairings have been developed by Lettovský et al. (2000) and Stojković et al. (1998). We can integrate these techniques within SimAir.

Our model allows airlines to evaluate the performance of a crew schedule in operations. Because disruptions inevitably occur, a schedule's performance in a model of operations may be a better measure of its quality than its planned cost. One challenge is to find crew schedules that perform well in the model. By considering delay and disruption probability distributions, it may be possible to construct crew schedules that have a better expected operational performance than a solution found by the current state-of-the-art method. Preliminary research is encouraging. In our computational examples, the crew schedule found using Monte Carlo methods had lower expected costs than the crew schedule that was found using state-of-the-art methodology. This indicates that considering uncertainty may lead to better plans.

Fleet assignment models assign a fleet type to each leg so that the airline maximizes revenue. After the fleet assignment is done, the airline finds an aircraft rotation that maximizes additional connection revenue and satisfies maintenance constraints. However, the rotation may be very sensitive to disruptions. For example, if the rotation does not have any short cancellation cycles, then when a leg is cancelled, the airline may have to cancel several legs to maintain plane flow balance. If the fleet assignment and the aircraft rotation models required many short cycles, the rotation would likely be less sensitive to cancellations. Our stochastic model can be used to test fleet assignments and aircraft rotations from these models.

## Acknowledgments
This work was supported through grants from The SABRE Group and United Airlines and NSF grant DMI-0085723. ILOG, Incorporated provided the linear programming and integer programming software that was used to find the crew schedules. A grant from the Intel Corporation provided the hardware used to solve the crew-scheduling problems and run SimAir. The authors would like to thank Barry Smith and Ladislav Lettovský from The SABRE Group and Eric Gelman, Srini Ramaswamy, Shivaram Subramanian, and Patrice Yapo from United Airlines for providing data and making suggestions to improve SimAir, and they thank the referees of this manuscript for their insightful comments.

## References
Adams, M. 2000. For fliers, May has been the cruelest month this year. *U.S.A. Today* (June 6) B1.

Barnhart, C., N. Boland, L. Clarke, E. L. Johnson, G. L. Nemhauser. 1998. Flight string models for aircraft fleeting and routing. *Transportation Sci.* **32** 208–220.

Boland, N. L., L. W. Clarke, G. L. Nemhauser. 2000. The asymmetric traveling salesman problem with replenishment arcs. *Eur. J. Oper. Res.* **123** 408–427.

*Bureau of Transportation Statistics*. 1998. On-time statistics. Available from ⟨www.bts.gov/ntda/oai/desc.html⟩.

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

Chu, H. D., E. Gelman, E. L. Johnson. 1997. Solving large scale crew scheduling problems. *Eur. J. Oper. Res.* **97** 60–268.

Clarke, L. W., E. L. Johnson, G. L. Nemhauser, Z. Zhu. 1997. The aircraft rotation problem. *Ann. Oper. Res.: Math. Indus. Systems II* **69** 33–46.

Dobbyn, T. 2000. U.S. prepares plan to cut summer flight delays. *Reuters Wire* (March 3).

Federal Aviation Administration. 1999. Federal aviation regulations. Available from URL: ⟨www.faa.gov/avr/AFS/FARS/far-121.txt⟩.

Gamache, M., F. Soumis, D. Villeneuve, J. Desrosiers, E. Gelinas. 1998. The preferential bidding system at *Air Canada*. *Transportation Sci.* **32** 246–255.

Gu, Z., E. L. Johnson, G. L. Nemhauser, Y. Wang. 1994. Some properties of the fleet assignment problem. *Oper. Res. Lett.* **15** 59–71.

Haeme, R. A., J. L. Huttinger, R. W. Shore. 1988. Airline performance modeling to support schedule development: An application case study. *Proc. Winter Simulation Conf.* 800–806.

Hane, C. A., C. Barnhart, E. L. Johnson, R. E. Marsten, G. L. Nemhauser, G. C. Sigismondi. 1995. The fleet assignment problem: Solving a large 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.

Klabjan, D., K. Schwan. 1999. Airline crew pairing generation in parallel. Technical report TLI/LEC-99-09, Georgia Institute of Technology, Atlanta, GA.

——, E. L. Johnson, G. L. Nemhauser. 1999a. Airline crew scheduling with regularity. Technical report TLI/LEC-99-11, Georgia Institute of Technology, Atlanta, GA.

——, ——, ——. 1999b. Solving large airline crew scheduling problems: Random pairing generation and strong branching. Technical report TLI/LEC-99-12, Georgia Institute of Technology, Atlanta, GA.

Kniker, T., 1998. Itinerary-base airline fleet assignment. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA.

Lettovský, L., 1997. Airline operations recovery: An optimization approach. Ph.D. thesis, Georgia Institute of Technology, Atlanta, GA.

——, E. L. Johnson, G. L. Nemhauser. 2000. Airline crew recovery. *Transportation Sci.* **34** 337–348.

*Olympic Airways*. 2001. *Olympic Airways'* fleet. Available from ⟨www.olympic-airways.gr⟩.

Phillips, D., N. Irwin. 2000. Flying into a storm of delays. *The Washington Post* (July 17) A01.

*Reuters Wire*. 2000. Researchers warn airline service declining. *Reuters Wire* (April 10).

Rosenberger, J. M., E. L. Johnson, G. L. Nemhauser. 2001. Rerouting aircraft for airline recovery. Technical report TLI/LEC-01-04, Georgia Institute of Technology, Atlanta, GA.

——, A. J. Schaefer, D. Goldsman, E. L. Johnson, A. J. Kleywegt, G. L. Nemhauser. 2000. A stochastic model of airline operations. Technical report TLI/LEC-00-06b, Georgia Institute of Technology, Atlanta, GA.

Schaefer, A. J. 2000. Airline crew scheduling under uncertainty. Ph.D. thesis, Georgia Institute of Technology, Atlanta, GA.

——, E. L. Johnson, A. J. Kleywegt, G. L. Nemhauser. 2001. Airline crew scheduling under uncertainty. Technical report TLI/LEC-01-01, Georgia Institute of Technology, Atlanta, GA.

Stojkovic, M., F. Soumis, J. Desrosiers. 1998. The operational airline crew scheduling problem. *Transportation Sci.* **32** 232–245.

Vance, P. H., A. Atamturk, C. Barnhart, E. Gelman, E. L. Johnson, A. Krishna, D. Madidhara, G. L. Nemhauser, R. Rebello. 1997. A heuristic branch-and-price approach for the airline crew pairing problem. Technical report TLI/LEC-97-06, Georgia Institute of Technology, Atlanta, GA.

Yang, W., Y. Zhu, Q. Tu, Y. Sheng. 1991. Simulation of commercial-aircraft reliability. *Proc. Annual Reliability and Maintainability Symposium* 112–119.

Yau, C.. 1991. An interactive decision support system for airline planning. *IEEE Trans. Systems, Man, and Cybernetics* **23** 1617–1625.

Yu, G., ed. 1997. *Operations Research in the Airline Industry*. Kluwer Academic Publishing, Boston, MA.

*Received: September 2000; revisions received: August 2001; accepted: August 2001.*