Coldstart: Fleet Assignment at Delta Air Lines

RADHIKA SUBRAMANIAN — *Delta Air Lines, Inc.*
1030 Delta Boulevard
Atlanta, Georgia 30320

RICHARD P. SCHEFF, JR. — *Delta Air Lines, Inc.*

JOHN D. QUILLINAN — *Delta Air Lines, Inc.*

D. STEVE WIPER — *Delta Air Lines, Inc.*

ROY E. MARSTEN — *Georgia Institute of Technology*
Atlanta, Georgia 30332

Delta Air Lines flies over 2,500 domestic flight legs every day, using about 450 aircraft from 10 different fleets. The fleet assignment problem is to match aircraft to flight legs so that seats are filled with paying passengers. Recent advances in mathematical programming algorithms and computer hardware make it possible to solve optimization problems of this scope for the first time. Delta is the first airline to solve to completion one of the largest and most difficult problems in this industry. Use of the Coldstart model is expected to save Delta Air Lines $300 million over the next three years.

Delta Air Lines has over 2,500 domestic flight departures every day. This includes flights to Canada and Mexico but excludes the other international routes. Delta has about 450 aircraft available to fly these flights. The aircraft are divided by aircraft type into *fleets* or groups, of which Delta has 10. The intricate pattern that the Delta aircraft fly along the route system is called the *schedule*. The schedule is the very heartbeat of an airline.

It has been said that an airline seat is the most perishable commodity in the world. Each time an airliner takes off with an empty seat, a revenue opportunity is lost forever. So the schedule must be designed to capture as much business as possible, maximizing revenues with as little direct operating cost as possible.

An airline combines the worst of both worlds. The capital-intensive quality of a manufacturing environment is combined

Copyright © 1994, The Institute of Management Sciences

0091-2102/94/2401/0104$01.25

TRANSPORTATION—AIR
PROGRAMMING, LINEAR—APPLICATIONS

INTERFACES **24:** 1 January–February 1994 (pp. 104–120)

with the low-profit environment of retail sales. Airlines are capital, fuel, and labor intensive. Survival and success depend on the ability to operate flights along the schedule as efficiently as possible. The profit-and-loss curve in the airline business has historically followed the peaks and valleys of the general economy. In this dynamic environment, where profits have historically been low and costs historically high, how well an airline plans and implements its schedule can very well determine its future. A small adjustment in the schedule may result in millions of dollars of additional revenue or in millions of dollars of losses. In planning schedules, continuous refinement is not a luxury; it is a requirement.

Both the size of the fleet and the number of different types of aircraft have an exponential impact on schedule planning. Fleets of different types of aircraft make up the total Delta fleet. We use the phrase *fleeting the schedule* to express assigning a particular set of aircraft to a particular set of markets.

The Coldstart project addressed the problem of fleeting the schedule. The basic trade-off is that if the airline uses too small a plane, it will leave potential passengers behind, while if it uses too large a plane, it will suffer the greater expense of the larger plane to transport empty seats. The goal is to have the right plane in the right place at the right time, but the many constraints on the way that planes can actually be operated make this difficult to accomplish.

The Coldstart model is a large-scale mixed-integer linear program that assigns fleet types to flight legs so as to minimize a combination of operating and passenger "spill" costs, subject to a variety of operational constraints, the most important of which is the number of aircraft available in each fleet. The model deals with a single day, which is assumed to be part of a repeating cyclic schedule. In practice, there are exceptions to the daily schedule, particularly on the weekends. At present, schedulers handle these exceptions manually. We are currently building an extension of the Coldstart model to handle the exceptions.

Our model assigns fleet types, not individual aircraft tail numbers, to the flight legs. Actual aircraft are routed after the model is solved to ensure that the solution is operational. Because of the hub-and-spoke nature of operations and large fleet sizes, it is always possible to obtain a feasible tail routing from the assignment recommended by the model.

The name Coldstart was inspired by an earlier schedule-planning tool, known informally as Warmstart, that was used by Delta Air Lines. Warmstart took a fleeted schedule that had been produced by the planners and tried to improve it using local swap heuristics. It could not, however, move very far away from its starting point, and this left the schedulers with most of the burden of coming up with a good assignment manually. The most serious drawback of the Warmstart system was that its improvements were so local that a poor input schedule would result in a poor output. Coldstart, being an optimization model, does not require an initial fleeting.

Until very recently, optimizing the fleet assignment for an airline as large as Delta would not have been possible. Today, however, improvements in mathematical

programming algorithms and computer hardware make it possible to solve optimization problems of this scope for the first time. Delta is the first airline to take advantage of this new opportunity.

Modeling the Fleet Assignment

The fundamental mathematical structure upon which the fleet assignment model is built is a time-space network. This network has a time line for each aircraft fleet at each city. This time line is circular and represents a repeating 24-hour period. Along a given time line, for example, the one for 757 in Boston, a node represents each point in time at which at least one event occurs. An event is the arrival or departure of a flight. Segments of the time line that connect these nodes are referred to as *ground arcs.* A *flight leg* is a single hop, that is, one takeoff and one landing. A flight leg is represented in the network by a set of *sky arcs,* one for each fleet that could be assigned to fly that leg. For example, assigning a 757 to the flight leg that leaves Atlanta at 6:21 AM and arrives in Boston at 8:45 AM is represented by a sky arc beginning at the 6:21 AM node on the (757, ATL) time line and ending at the 8:45 AM node on the (757, BOS) time line. Hence the time lines for a given fleet at different cities are connected by sky arcs; however, the time lines for different fleets are not connected. Figure 1 provides a simple illustration of how the time lines for the 767-200 fleet at Atlanta (767, ATL) and San Juan (767, SJU) are connected.

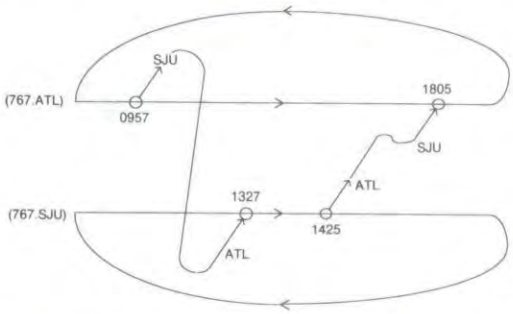

Figure 1: Time line network for a very simple airline with only one fleet (767) that serves only two cities, Atlanta (ATL) and San Juan (SJU). Each arrival and departure is labeled with a time (24-hour clock).

These (fleet, city) time lines allow us to express conservation of flow equations for the aircraft. For each sky arc, we define a binary assignment variable that takes the value one if that fleet is assigned to that leg or the value zero otherwise. For each ground arc, we define an integer variable that counts the number of planes of that type on the ground at that city during that time interval. At each node, there is a conservation equation. For example, for the 6:21 AM node on the (757, ATL) time line, the equation says

> (number of 757s on the ground in ATL just before 6:21 AM) − (number of 757s that depart from ATL at 6:21 AM) + (number of 757s that arrive at ATL at 6:21 AM) = (number of 757s on the ground in ATL just after 6:21 AM).

The equations are presented in detail in the appendix. For a development of a similar model, see Berge and Hopperstad [1993] and Hane et al. [1993].

The conservation equations make up the bulk of the model constraints. There is one constraint for each fleet at each node, where each flight results in two nodes. Hence, if there are 10 fleets and 2,500 flights per day, and if every fleet could fly every leg (which is not true due to operational limitations), and if all departure and arrival times for a given fleet at a given

city were unique (also not true), then we would have $2,500 \times 2 \times 10 = 50,000$ conservation equations. The actual number for Delta's daily model is closer to 30,000.

In the Coldstart model, note that a plane arriving at a city can be paired with any later departure. This formulation is far more powerful and realistic than the fleet assignment model of Abara [1989]. In the model discussed by Abara, each feasible turn and aircraft combination represents a decision variable. To limit the size of his model, each arrival could be paired with at most the next five departures. Abara's formulation is further limited in that it does not allow the use of the model to be extended to other schedule-related applications like fleet planning and route development. This ability to pair an arrival with any later departure is very important at the hubs, where there can be as many as 65 arrivals and 65 departures at a complex. (A *complex* is a set of arrivals and departures that connect to each other, a typical feature of the hub-and-spoke operation.)

For a given flight leg, as many as 10 different binary variables represent assigning that leg to the 10 different fleets. These assignments are mutually exclusive, so there is a multiple-choice constraint that says that the sum of these zero/one variables must be exactly one. There is one such cover constraint for each of the 2,500 flight legs. These cover constraints are expressed as "special ordered sets, type 3" [Druckerman, Silverman, and Viaropulos 1991], which considerably reduces the time for solution of the mixed-integer program.

Certain pairs of flight legs must be assigned to the same fleet to provide one-stop, through service. For example, the plane that flies the 1:41 PM flight from Atlanta to Boston, arriving at 4:10 PM, must also fly the 5:25 PM flight from Boston to Montreal. Thus these two legs must be assigned to the same fleet. This is easily modeled by an equation of the form "$X_1 - X_2 = 0$" for each fleet. The typical daily model involves about 250 such required hookups, giving rise to about 1,500 constraints.

The number of planes in each fleet, which is the scarce resource in the assign-

An airline seat is the most perishable commodity in the world.

ment problem, is obviously limited. To capture this in the model, we build a size constraint for each fleet. The size constraint for 757s includes every 757 that is in the air at midnight Atlanta time and also counts the 757s on the ground at midnight Atlanta time. To prevent getting an infeasible solution, we allow the model to use extra planes but at a very high cost. This is preferable to getting an infeasibility that may be hard to diagnose.

Aircraft need to be maintained at periodic intervals. Short maintenance can be done while a plane is sitting on the ground during the day. Longer maintenance must be performed at night. Part of the input to the Coldstart model is a list of overnight maintenance requirements. For example, one 757 must have a 12-hour maintenance check every night. For each such requirement, there is a list of maintenance bases where the check can be done. The model must then insure that a 757 is available in

one of these maintenance bases for 12 hours. To model this, we introduced a new set of maintenance arcs. A *maintenance arc* represents a maintenance opportunity. It begins at an evening arrival node at a maintenance city and ends at a morning departure node that is at least 12 hours later (for a 12-hour check). There are several such opportunities for each requirement, perhaps at several alternative cities. There is then a multiple choice constraint over the opportunities that correspond to each requirement. A typical model has about 30 maintenance requirements. An important modeling extension to the major maintenance requirement is that the model can select the best city at which to perform a certain maintenance. This feature is very useful in analysis of the schedule and in the planning mode.

The individual fleets are grouped into *pilot aggregates* that can be flown by the same pilots. For example, the Boeing 757, 767-200, and 767-300 fleets can all be flown by the same pilots. For each pilot aggregate, there is a limit on the number of flying hours per day that can be assigned to that set of pilots. Each assignment variable must appear in the flying hours constraint for the pilot aggregate that it belongs to, with a coefficient equal to the number of flying hours for that leg. These constraints, like the cover constraints, have the effect of coupling the fleets together.

Airlines are capital, fuel, and labor intensive.

There are other crew-related considerations that must be built into the model. Just as planes need maintenance, pilots need rest. A fleet assignment that is good from the point of view of aircraft scheduling may be very bad from the point of view of crew scheduling. For example, suppose there is only one 757 flight in and out of Boise each day and no flights by any other fleet in the same pilot aggregate. If this flight arrives at 11:00 PM and leaves at 7:00 AM, and if the arriving crew has to have an overnight rest break of at least 10½ hours, then this crew will have to remain in Boise until 7:00 AM the day after—32 hours after its arrival. This is very expensive in terms of crew costs. What this means for fleet assignment is that it is not good to have too few flights by a pilot aggregate into a city, or, equivalently, to have too many different pilot aggregates serving a city. In particular, late arrivals should be paired with midday departures so that a crew can fly out of the city soon after their rest break. To model this, we add constraints that result in an objective function penalty if the number of midday departures by a pilot aggregate is less than the number of late arrivals by the same pilot aggregate. We also try to find "10:30 opportunities" (10½ hours are the required time for a legal crew break). A 10:30 opportunity is a pairing of an evening or late night arrival by a certain pilot aggregate and a departure by the same aggregate at least 10½ hours later, such that there is at least one plane of that aggregate on the ground at that city for the whole 10½+ hour period. We find these legal crew break opportunities by building crew time lines [Johnson 1992] for fleet aggregates at certain busy cities. A crew arriving at an airport can leave from a coterminal after a

legal break. For example, LaGuardia, John F. Kennedy, and Newark airports form the New York area coterminal. The crew constraints deal with coterminal sets of cities.

The operational restrictions are coded into the input data for Coldstart. Certain aircraft types are operationally constrained from serving specific flight segments or legs in the Delta schedule. These restrictions are due to aircraft performance limits, such as takeoff or landing weight limits at specific airports, or the range capability of the aircraft. The range capability of the aircraft is determined by its fuel capacity, the altitudes of the airports being served, and the distance between two airports. Not all aircraft are equipped with overwater navigational equipment, life rafts, and so forth; segments requiring such equipment must be covered by one of the specific fleet types with this equipment.

Other restrictions coded into the input data include airport restrictions, such as stage 3 noise restrictions, and general arrival and departure curfews. Aircraft are certified to meet various noise standards, such as stage 2 or 3, in accordance with Federal Aviation Rules and International Civil Aviation Organization standards. Noise certification standards are complex; they involve a three-point evaluation scheme (that includes noise levels measured at takeoff, sideline, and approach), computation of an effective perceived noise level, and noise limits based on maximum takeoff weight. A stage 3 certified aircraft is generally quieter than a stage 2 certified aircraft of the same weight. Airports, such as Seattle, Washington, and San Francisco, restrict departures and arrivals between 11 o'clock in the evening and 6 o'clock in the morning to aircraft meeting stage 3 noise requirements. Orange County, California and Washington, DC's National airports are airports that allow only certain aircraft types at all. Orange County, California allows only certain aircraft types, which include 737-300, 737-400, and 757, to arrive or depart during the airport's operations; Washington, DC's National airport allows only 757 departures and MD88 arrivals between 10 o'clock in the evening and 6:49 in the morning.

While some noise restrictions are accounted for in the input data, noise restrictions that are enforced as a percentage of landings and takeoffs that can be stage 2 aircraft are modeled by a blending constraint. In San Diego, California, for example, at most 25 percent of the departures may be stage 2 aircraft.

A very serious modeling difficulty that we overcame has to do with the arrival time of each flight leg. Using the time-space network to balance the flow of aircraft means that a plane that arrives at 10:50 AM, for instance, can leave on any later departure. This implies that by arrival time we do not mean the time when the plane lands on the runway, or even the time when it pulls up to the gate. What we really mean is the time when it is ready to go out again. The time-space network has to be built with ready times rather than with scheduled arrival times. The time that an aircraft takes to be ready for the next

We spent an immense amount of effort fine tuning and cleaning our data.

takeoff depends on whether it is continuing with the same flight number or is changing flight numbers. For a continuing flight, we assume that the aircraft needs less time to get ready for the next takeoff. The difference in these ready times is only about five to 10 minutes, but since the schedule is so tight, this difference makes a big impact on the overall fleeting. The planners, at their discretion, will sometimes allow an aircraft ready time to go below the minimum if this means a good connection somewhere further down in the schedule. This tightening of the schedule to allow for high revenue connections is a very important modeling feature. We have modeled this by adding flight arcs for missed connections with a cost proportional to the time shaved from the usual ready time. The variables corresponding to these arcs go into the cover constraint for that flight leg.

Another related issue that affects the arrival time by a few minutes is the fact that the different fleets fly at different speeds. We have grouped the fleets into speed classes (different from the grouping into pilot aggregates). Flying time is more complex than just distance divided by speed. For example, it depends on the direction, season, and fleet type. Another factor that affects arrival time is the taxi time, which varies by fleet type and airport. To model these effects, we build the sky arcs differently depending on nominal aircraft speed and make small adjustments to the scheduled arrival and departure times to avoid violating the arrival and departure banks of complexes at the hub cities. The model generator retimes the flights to avoid infeasible connections and to make feasible connections based on the various aircraft speeds. Hence the resulting assignment is feasible from an operational standpoint.

Last, we modeled pilot training as a constraint. Pilot training requires the availability of specific aircraft types at specific cities for fixed periods of time daily.

The Objective Function

We have used three objective functions in Coldstart thus far. The primary objective has been cost minimization. In the cost minimization objective, the goal is to minimize the sum of the operating cost, spill cost, and any applicable penalties. The operating cost by fleet type is extracted from the accounting ledger for each leg. The operating cost consists of several components that can vary by fleet type and includes crew (pilot and flight attendant) cost, fuel cost, landing fees, and a maintenance burden. Spill cost is estimated, and depends upon the demand for a leg, aircraft capacity, recapture rate, and the revenue of the lost passenger.

Because some constraints in the model are not hard and fast but rather soft, we incorporated a set of bonuses and penalties in the objective function to help make the solutions operationally feasible. For example, a penalty exists for flying a wide-body fleet into a city that is not currently served by wide-bodies, and a bonus exists for flying a fleet into a city that is a crew base for its pilot aggregate.

Spill is the number of passengers not carried because aircraft capacity is insufficient. It is the unconstrained demand lost due to insufficient aircraft capacity on a leg [Swan 1983, 1992a, 1992b]. Spill is caused by the truncation of the demand distribution beyond the aircraft capacity (Figure 2).

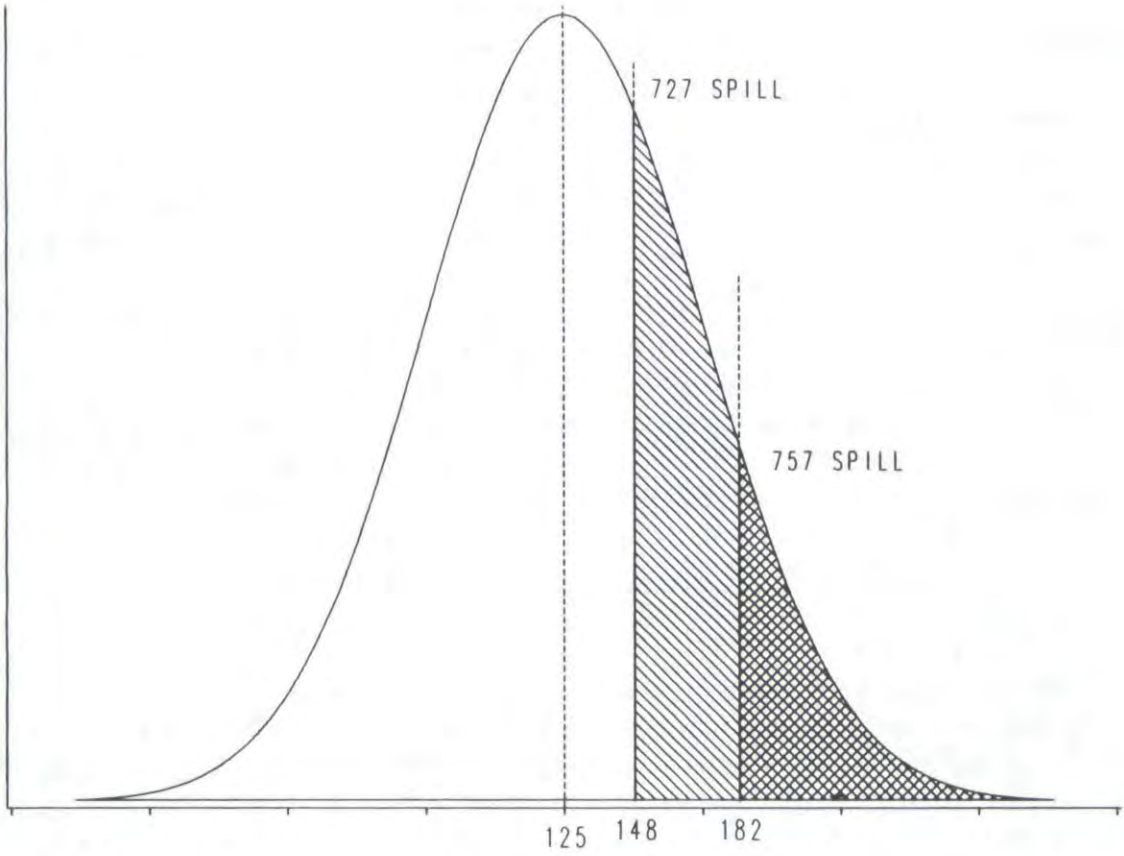

Figure 2: The figure shows a normally distributed passenger demand with a mean of 125 and standard deviation of 45. The area under the demand curve to the right of a vertical line at the capacity of an aircraft represents full flights with unsatisfied demand. Spill for a Boeing 727 with a capacity of 148 and spill for a Boeing 757 with a capacity of 182 are represented by the areas to the right of the vertical lines at 148 and 182, respectively.

The variation in daily passenger demand can be attributed to differences in day-of-the-week demand, seasonality, cyclic effects, and random variation. High variability can result in large passenger spills. After finding the unconstrained mean and standard deviation, one can calculate the expected spill for any size plane.

Spilled passengers are either recaptured on other Delta flights or lost to competitors. Spill becomes a cost to an airline when spilled passengers are lost to competitors or other modes of transportation. Using a market-share model, we estimate the percentage of spilled passengers that are recaptured on other Delta flights on a leg-by-leg basis. This percentage is called the recapture rate. The loss rate for spill is simply one less the recapture rate.

The product of the expected spill and the loss rate for spill gives the lost spilled passengers. We convert this into dollars by multiplying it by an estimate of the average revenue per spilled passenger on that leg. The resulting lost spilled revenue is the spill cost.

The spill cost and any applicable penalties and bonuses (which are negative penalties) are added to the operating cost to estimate the total cost of operating a particular fleet type on a leg. In addition to minimizing the cost, we have the flexibility to run the model for other objective functions. For fleet-planning purposes, we can change the objective function to minimize the total number of planes used to fly the schedule. We introduced a third objective function for route-development purposes. If we change the objective function from cost minimization to profit maximization, we can modify the cover constraints to determine when to add new service or drop existing flights.

Solution Technique

The typical size of the daily Coldstart model is about 40,000 constraints and 60,000 variables. There are roughly 20,000 binary variables (assignment of fleets to legs and maintenance variables) and 40,000 general integer variables. Hence we have a large mixed-integer programming problem to solve. Our solution strategy is to use the OB1 interior point code [Lustig, Marsten, and Shanno 1991, 1992] to solve the problem as a linear program; fix some or all of the binary variables that are at 1.0 in the LP solution; use these fixed variables to reduce the size of the problem; and solve the resulting smaller mixed-integer problem with the OSL mixed integer programming code [Druckerman, Silverman, and Viaropulos 1991].

Before attempting to solve the linear program, OB1 uses general algebraic reduction techniques to reduce the problem size. First it uses the "lonely-plus/lonely-minus" reduction [Lustig and Marsten 1993]. Lonely-plus works as follows. Suppose we have an equation with just one positive coefficient, several negative coefficients, and a nonnegative right-hand side. Suppose further that all of the variables must be nonnegative. Then the variable with the lonely-plus can never be negative, and we can use this equation to substitute it out of the model. Thus the model is reduced by one equation and one variable. The reductions obtained in this way include some very natural node aggregations that could be done at the matrix-generation stage. Suppose that on a (fleet, city) time line we find a sequence of nodes that represent only arrivals, followed by a sequence of nodes that represent only departures. Then any of these arrivals could connect to any of these departures, and all of these nodes can be aggregated into a single node. Figure 3 illustrates the concept of node aggregation. This eliminates several nodes and hence several balance equations and several integer variables for the

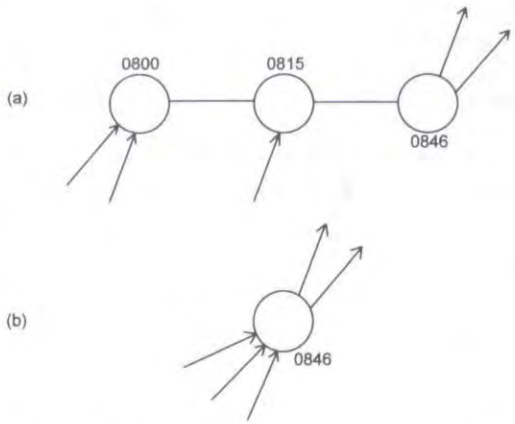

Figure 3: Node aggregation in the time line network: (a) shows the actual network; (b) shows a mathematically equivalent network. A plane that arrives at 0800 or 0815 can be sent out on either of the 0846 departures.

intervening ground arcs. The lonely-plus/lonely-minus reduction also finds and eliminates the required hookup constraints, which are of the form, $X_1 - X_2 = 0$. This has the effect of combining two separate variables into a single variable. Other reductions that are found do not have a simple interpretation in terms of the network.

The second model reduction used by OB1 is to eliminate dependent rows. In using the interior-point method, we do not introduce artificial variables for the equality constraints. This means that there may be dependent rows, and in fact there are quite a few among the conservation-of-flow constraints. We detect these by performing a Gaussian reduction on the coefficient matrix. Finally, OB1 uses the standard sort of model reductions described by Brearley, Mitra, and Williams [1975]. The model to be solved ends up being about 12,000 rows and 30,000 variables. (For example, the December 1992 schedule resulted in a model with 41,325 rows and 62,088 variables and was reduced to 12,811 rows and 33,475 variables.)

Computational experiments described in Hane et al. [1993] show convincingly that the interior-point method dominates the simplex method for this class of problems. For example OB1 uses the interior-point method to solve the December 1992 schedule in 45 iterations, taking 43 minutes on an IBM RS/6000 (Model 530) work station. The OSL primal-simplex code solves the same model in 356,854 iterations, taking 19 hours on the same work station. OB1 uses the predictor-corrector version of the primal-dual interior-point method [Lustig, Marsten, and Shanno 1992]. (The OSL predictor-corrector, primal-dual barrier code [Druckerman, Silverman, and Viaropulos 1991] is very similar to OB1 but does not perform as well on the fleet assignment models because the OSL preprocessor does not remove dependent rows and does not do the lonely-plus/lonely-minus reduction if the variable being substituted out has more than one other nonzero coefficient.)

The LP solution of the daily model typically assigns a unique fleet to about 80 percent of the flight legs, or about 2,000 of the 2,500 legs. This leaves about 500 legs with two or more fleets assigned at a fractional level. If all of the binary variables that are at one in the LP solution are fixed, then the same kind of algebraic model reductions discussed above will reduce the problem down to about 4,000 rows and 6,000 variables. Near optimal solutions to models of this size can be found within an hour or two of branch-and-bound search by the OSL MIP code [Druckerman, Silverman, and Viaropulos 1991]. By near-optimal, we mean within 0.1 percent of the LP objective value of the full model. Of course, the reduced problem may be infeasible because of the fixed variables. We have developed heuristics to fix some but not all of the X-variables that are at one in the LP solution so as to prevent infeasibility. Fortunately, infeasibility is rarely a problem when we are solving the large all-fleet models—so rare that we have not encountered it. Infeasibility occurs more often when variables are fixed for two or three fleet models, and in this case, OSL can solve the whole problem from scratch without help from OB1.

Implementation and Operational Impact

We developed Coldstart to take advan-

tage of the economic analysis that had been done for an earlier system that made small fleeting changes to improve profit. The earlier system, Warmstart, was essentially a local swapper and would look for interchangeable paths of two to four flights that could be switched to improve profit. The following example illustrates the type of swap that the system produced (all times are in Eastern Standard Time):

Atlanta to Birmingham 1100 1140,
Birmingham-Dallas/Ft. Worth
1230 1330, fleet: 72S;
Atlanta to Jackson 1100 1210,
Jackson-Dallas/Ft. Worth
1245 1330, fleet: 757.

Clearly these two paths could be interchanged if the end result was desirable, because the situation before and after the swap is identical. The rest of the schedule can be ignored as far as this swap is concerned.

While the scope of the changes that can be made with this procedure is much smaller and the applicability is limited, the economic analysis is essentially the same. The question of the cost of a particular type of aircraft on a particular segment must be identified and then used to improve the schedule.

Delta used the Warmstart system for over a year for fleeting decisions—and thus the economic assumptions and forecasts. Once it had decided to use this analysis for fleeting decisions, the leap to accepting Coldstart recommendations was a smaller one.

Once Delta started to use the local swapper, its shortcomings became apparent. These weaknesses spawned the effort to generate a more global and far-reaching fleeting system. Delta formed the operations research team initially to work on the Coldstart model.

Once Coldstart was operational, the decision to use it was very easy. The schedule planners analyzed the difference between their initial schedules and the recommended Coldstart fleetings. They felt that the solutions Coldstart produced were clearly superior. The planners can now determine analytically which flights need to be upgraded and which downgrades will be costly, and they can assess a schedule once the fleeting is done. The problem prior to Coldstart was making desirable fleetings operationally feasible.

The economic analysis necessary to run the fleet optimization makes estimating the benefits fairly simple. The cost of a fleeting is determined by summing the segment costs for the assigned fleets. The costs between fleetings can then be compared. In the meantime, the operating departments provide feedback on other unmodeled costs that may be affected by the changes, including crew, maintenance, and station costs.

Delta has two scheduling groups which work on each of the six to eight schedule periods in a year. One of the groups, the planning group, works on a schedule beginning eight to nine months before it starts and provides a tentative schedule to the operating departments seven months ahead of time. This schedule is used to set staffing levels at the pilot bases and to provide a good starting solution for the current schedules group. The other scheduling group, the current group, starts work on the final version of the schedule about four months early using the planning version as

a starting point. While the group may make changes, they may not exceed the pilot staffing levels at any of the crew bases that were set as a result of the planning schedule. Since they work several months closer to the actual schedule date, the current group can take advantage of competitive changes, growing or declining markets, or other changes in strategy to refleet to better advantage. In addition, there are generally some changes to the actual legs in the schedule between the time the planning group issues a schedule and the current group does. The available fleet count may change also.

Coldstart is the first OR application of this magnitude developed and implemented inside Delta.

Both the planning and current scheduling groups use the Coldstart model extensively, although the planning group generally has more opportunity to make sweeping changes. The current group, in addition to having to keep within the crew base constraints, must schedule all the weekend cancellations and other exceptions to the assumption of the repeating 24-hour cycle. Because of the reduced lead time, the current group generally will have a more accurate demand forecast; it will call for some refleetings from the initial schedule from planning.

Both groups tend to use Coldstart in a similar fashion. Initially, they include all the fleets in a run to get a good starting solution. Invariably, there are assignments that are undesirable for one of a number of reasons that cannot easily be captured in the model. This first run takes from one to three hours since it includes all 10 fleets and 2,500 flight segments. From this solution, the schedule planners can identify the problems and work them out individually or in groups. They generally solve these problems quickly as subproblems involving from two to five fleets, often in five minutes or less. This tuning of the initial solution will generally continue until the deadline for issuing the schedule.

Coldstart has changed the basic task of a schedule planner at Delta. While the model performs all the schedule changes, and very fast, the task of the planner involves more analysis of these changes. Planners spend a lot more effort in analyzing the various cost numbers that are used to drive the model. With this tool, they can test various scenarios in a short amount of time and choose the best one. As a result, the flight schedule at Delta has changed considerably since the model went into effect.

Because the old method of fleeting is no longer used, benchmarks are not available. While the question, "how much better is this Coldstart fleeting than what we would have achieved manually," was important initially; a better question now is, "how much better can a Coldstart solution become?" We will always be changing the model, and we can use an infinite number of parameter settings; this is the key to future enhancement. Coldstart allows the user to produce two different schedules using different constraints or parameters and allows him or her to answer these types of questions with the same comparative procedure.

Planners can use Coldstart in a number of ways. Initially, they can use it to get a fleeting that is close to the desired one. They accomplish this by doing one or more runs using all the fleets. Once that is accomplished, they work on segments that need to be refleeted because of factors that cannot be modeled. Coldstart can be used to solve smaller subproblems—usually the best and fastest way to address individual problems. For example, a planner may wish to upgrade a flight from Atlanta to Dallas/Fort Worth from a 76S with 254 seats to an L10 with 302 seats. The planner can pull the 76S and L10 fleets out as if they were a separate airline, target the desired upgrade and see what refleetings are necessary to accomplish the change. This can often be done in just a minute or so, depending on the size of the subproblem. A desirable solution to a targeted flight can be accepted and put into the schedule base if desired. Because the system is so flexible, it can be used to solve small fast subproblems and thus accomplish any number of desired goals.

The schedules department sends out schedule proposals to the operating departments and then gets responses specifying problems and requesting changes. It takes these responses and revises the schedule to best address the problems and the needs of the operating departments. The targeted subproblem approach with minimal changes works well in the late stages of the schedule development process. Once schedules has issued schedule proposals, it must keep the number of changes before the schedule is finalized to a minimum. This keeps the number of refleetings small while accomplishing the necessary changes.

Extensions

The Coldstart model has grown in three different directions beyond its original purpose. These are fleet planning, route development, and transition planning. Implementing these extensions was easy because of the way the initial model was built. Fleet planning is a natural extension of fleet assignment. Instead of treating the fleet sizes as fixed, it is easy to model the possibilities of acquiring additional aircraft or retiring existing aircraft. In addition to operating cost, some measure of the ownership cost must be included in the model. The success of the model as a fleet planning tool depends mainly on the objective cost figures. At Delta we have spent an immense amount of effort fine tuning and cleaning our data for these runs.

If the objective of the model is changed from cost minimization to profit maximization, then the optimizer can be allowed to choose which legs to fly, rather than being required to fly all legs. This means that the model can be used for developing routes by considering the addition of new legs or the deletion of existing legs.

A transition problem arises whenever a schedule change occurs, major or minor. For example, the schedules changed on Sunday, April 4, 1993 when daylight saving time began. On the night of April 3, the planes were all in their correct positions for continuing the winter schedule. But these were not the correct positions for beginning the spring schedule! One solution would be to ferry planes, empty, in the dead of night, to their correct positions. Instead, Delta alters the assignment of fleets to legs during the last day of the old

schedule and the first day of the new schedule, sometimes having to cancel a few flights. This is a very time-consuming, manual planning job. We have developed and implemented a separate version of the Coldstart model to solve this problem. The objective function tries to minimize the number of fleeting changes required, with preference given to changes involving the same pilot aggregate or seating capacity.

The model has been built so that the user can select any subset of constraints that are applicable for his or her run. We have taken extensive effort to ensure that for a given run the model is never infeasible. This includes infeasibility due to fleet size, turn times, maintenance requirements, noise constraints, and pilot block hours.

Financial Impact of the Coldstart Project

When we began the Coldstart project in September 1991, we were not sure that we would be able to model and solve such a very large and complex system. The following factors were essential to the success of this project. First, we had someone with experience in the Delta fleet scheduling department working full time as part of the development team. He made sure that the model produced operationally feasible solutions. Second was the effort devoted to modeling the details of the system, particularly the turn time versus through time issue, including all of the many complications surrounding almost missed connections and aircraft speed. Third was state-of-the-art algorithms and software. OB1 and OSL are both cutting-edge codes, and OB1 was improved to meet the computational challenge posed by these models. Finally, the availability of powerful and inexpensive work stations put the equivalent of three completely dedicated mainframes at our disposal for development and testing.

Since the September 11, 1992 schedule, which was the first schedule that used Coldstart extensively, the model has been used for all of Delta's schedules. The refleetings from the September 11, 1992 schedule saved an estimated $55,000 per day over the schedule that would have been used. The planning group first used Coldstart for the December 15, 1992 schedule and obtained an estimated savings of more than $100,000 per day. All other schedules have been worked exclusively with Coldstart, which eliminates the benchmark comparison to the schedule as it would have been developed under the previous methodology. However, with additional experience and confidence, Delta has used the model much more heavily in developing subsequent schedules. The savings in the June 1, 1993 to August 31, 1993 schedule have been estimated at $220,000 per day. The savings from the model have increased from schedule to schedule as the planners gain more confidence and accept more and more of the model's recommendations. The success of this process has been mainly due to the willingness at Delta to understand and implement the model, sometimes even changing the way operations have traditionally been done at Delta.

A substantial percentage of the cost savings results from reduced direct operating cost. This amount is much more easily tracked than the reduced spilled revenue. It is a straightforward process to estimate the direct operating cost of two fleetings

over the same set of legs. Even here, there are some pitfalls. For example, a fleeting that allows you to save pilot costs and reduce the pilot head count by 100 will not actually save that money in the short term unless you lay off pilots, which Delta has not done. In the long run, the savings will be realized as attrition occurs or as Delta is able to increase its level of service without adding pilots.

Even when a schedule has been flown, it is very difficult to estimate how much more or less revenue would have been generated with a different fleeting. There is no record of spill or recapture even after it has occurred. Estimating this part of the objective function from four to eight months out is difficult. In addition to estimating the unconstrained demand, one must estimate the unconstrained variance for each flight. Since most flights have truncation in the historical data, statistical analysis must be used to estimate these parameters. We have used the techniques described by Swan [1983, 1992a, 1992b] in conjunction with techniques developed by Delta to estimate these key values. To this point, Delta has been pleased with both the cost savings and the revenue generation from the model.

Coldstart is the first operations research application of this magnitude that has been developed and implemented inside Delta Air Lines. Its success ensures that Delta will be an eager user of operations research techniques in the future.

APPENDIX

We present here an algebraic formulation of the basic model, which does not include maintenance, pilot training, pilot hours, crew breakout, crew 10:30 rest, and noise constraints.

The Basic Fleet Assignment Model

Sets:

CITY—a set of cities, indexed by i,

FLEET—a set of aircraft types, indexed by k,

LEG—a set of flight legs, indexed by l, and

TIME—a set of times, {0000, 0001, . . . , 2359}, indexed by t.

Parameters:

$cost(k,l)$ = cost if an aircraft of fleet k flies leg l (∞ if fleet k cannot fly leg l),

$turn(k,i)$ = turn time for fleet k in city i; the minimum time required between the arrival and subsequent departure of the same aircraft,

$size(k)$ = number of aircraft available in fleet k,

$orig(l)$ = origin city of leg l, an element of CITY,

$dest(l)$ = destination city of leg l, an element of CITY,

$depart(l)$ = departure time of leg l, an element of TIME, and

$arrive(l)$ = arrival time of leg l, an element of TIME.

Derived Sets:

NODES(k,i) = $\{t \in \text{TIME} \mid t = depart(l)$ for some $l \in$ LEG such that orig$(l) = i$, or $t = arrive(l)+turn(k,i)$ for some $l \in$ LEG such that $dest(l) = \text{i}\}$.

NODES(k,i) is the set of times when an arrival or departure of an aircraft of type k may happen at city i. For any $t \in$ NODES(k,i), we use $t+$ to denote the next time, and $t-$ to denote the previous time, with the circular assumption that 2359+ = 0000 and 0000− = 2359.

INTO(k,i,t) = $\{l \in \text{LEG} \mid dest(l) = i, arrive(l)+turn(k,i) = t\}$, $\forall k \in$ FLEET, $i \in$ CITY, and $t \in$ NODES(k,i).

OUTOF(k,i,t) = $\{l \in \text{LEG} \mid orig(l) = i, depart(l)=t\}$, $\forall k \in$ FLEET, $i \in$ CITY, and $t \in$ NODES(k,i).

$count_air(k)$ = $\{(k,l) \mid \text{cost}(k,l) < \infty$ and $arrive(l)+turn(k, dest(l)) < depart(l)\}$

$count_ground(k)$ = $\{(k,i,t) \mid t \in$ NODES(k,i) and $t+ < t\}$

$count_air(k)$ captures the set of legs where

an aircraft of fleet k may be in the air at midnight.

count_ground(k) captures the set of cities where an aircraft of fleet k may be sitting on the ground at midnight.

HOOK $\subseteq$ LEG $\times$ LEG

$(l_1,l_2) \in$ HOOK means that *dest*(l_1) = *orig*(l_2) and that the same fleet must be chosen for both l_1 and l_2.

Variables:

$X_{k,l}$ = 1 if fleet k is assigned to leg l;
= 0 otherwise.
$\forall k \in$ FLEET and $l \in$ LEG, such that cost(k,l) $< \infty$.

$Y_{k,i,t}$ = number of aircraft of fleet k on the ground at city i from time t to time $t+$.
$\forall k \in$ FLEET, $i \in$ CITY, and $t \in$ NODES(k,i).

Z_k = number of aircraft of fleet k that are used.

Constraints:

BALANCE(k,i,t):

$$\sum_{l \in INTO(k,i,t)} X_{k,i} - \sum_{l \in OUTOF(k,i,t)} X_{k,i} + Y_{k,i,t-} - Y_{k,i,t+} = 0,$$

$\forall k \in$ FLEET, $i \in$ CITY,

and $t \in$ NODES(k,i).

COVER(l):

$$\sum_{\{(k,l) \mid cost(k,l) < \infty\}} X_{k,l} = 1, \forall l \in \text{LEG}.$$

SIZE(k):

$$\sum_{(k,l) \in count_air(k)} X_{k,l} - \sum_{(k,i,t) \in count_ground(k)} Y_{k,i,t} - Z_k = 0,$$

$\forall k \in$ FLEET.

$0 \le Z_k \le size(k)$

HOOKUP(l_1,l_2,k):

$X_{k,l_1} - X_{k,l_2} = 0, \quad \forall k \in$ FLEET,

$(l_1,l_2) \in$ HOOK.

Objective:

Minimize COST = $\sum_{k \in \text{FLEET}, l \in \text{LEG}} \text{cost}(k,l) * X_{k,l}$.

Additional constraints can be added for maintenance requirements, crew considerations, pilot hours, pilot training, noise restrictions, and others.

References

Abara, J. 1989, "Applying integer linear programming to the fleet assignment problem," *Interfaces*, Vol. 19, No. 4, pp. 20–28.

Berge, M. E. and Hopperstad, C. A. 1993, "Demand driven dispatch: A method for dynamic aircraft capacity assignment, models and algorithms," *Operations Research*, Vol. 41, No. 1, pp. 153–168.

Brearly, A. L.; Mitra, G.; and Williams, H. P. 1975, "Analysis of mathematical programming problems prior to applying the simplex method," *Mathematical Programming*, Vol. 8, No. 1, pp. 54–83.

Druckerman, J.; Silverman, D.; and Viaropulos, K. 1991, "IBM optimization subroutine library, guide and reference, release 2," document number SC230519-02, IBM, Kingston, New York.

Hane, C. A.; Barnhart, C.; Johnson, E. L.; Marsten, R. E.; Nemhauser, G. L.; and Sigismondi, G. 1993, "The fleet assignment problem: Solving a large-scale integer program," working paper, School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia.

Johnson, E. L. 1992, IBM Corporation, personal communication.

Lustig, I. J. and Marsten, R. E. 1993, "Techniques for presolving linear programs," working paper.

Lustig, I. J.; Marsten, R. E.; and Shanno, D. F. 1991, "Computational experience with a primal-dual interior point method for linear programming," *Linear Algebra and Its Applications*, Vol. 152, pp. 191–222.

Lustig, I. J.; Marsten, R. E.; and Shanno, D. F. 1992, "On implementing Mehrotra's predictor-corrector interior point method for linear programming," *SIAM Journal of Optimization*, Vol. 2, No. 3, pp. 435–449.

Lustig, I. J.; Marsten, R. E.; and Shanno, D. F. forthcoming, "Interior point methods for linear programming: Computational state of the art," *ORSA Journal on Computing*.

Mehrotra, S. 1992, "On the implementation of a (primal-dual) interior point method," *SIAM*

Journal of Optimization, Vol. 2, No. 3, pp. 575–601.

Swan, W. M. 1983, "Traffic losses at high load factors," Proceedings of the 23rd Annual AGIFORS Symposium, Olive Branch, Mississippi.

Swan, W. M. 1992a, "The value of yield management information for routing and scheduling," presentation to AGIFORS Reservations and Yield Management Study Group.

Swan, W. M. 1992b, "Spill model revisited," working paper, Boeing Commercial Airplane Company, Seattle, Washington.

Copyright 1994, by INFORMS, all rights reserved. Copyright of Interfaces is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.