Applying Integer Linear Programming to the Fleet Assignment Problem

JEPH ABARA

American Airlines
Decision Technologies
PO Box 619616
Dallas/Fort Worth Airport, Texas 75261-9616

We formulated and solved the fleet assignment problem as an integer linear programming model, permitting assignment of two or more fleets to a flight schedule simultaneously. The objective function can take a variety of forms including profit maximization, cost minimization, and the optimal utilization of a particular fleet type. Several departments at American Airlines use the model to assist in fleet planning and schedule development. It will become one of 10 key decision modules for the next generation scheduling system currently being developed by American Airlines Decision Technologies.

American Airlines' schedule comprises a list of over 2,300 flights per day to over 150 different cities utilizing over 500 jet aircraft. This schedule is produced by considering an existing set of flights, traffic revenue forecasts, available resources such as aircraft, gates, and associated operating costs.

Within a schedule, there is a repeating pattern of flights, with the pattern covering one day or several days, usually a week. The goal of the fleet assignment process is to assign as many flight segments as possible in a schedule pattern to one or more aircraft types (American currently operates ten fleet types) while optimizing some objective and meeting various operational constraints.

The best aircraft for each flight leg is not always the one with the highest benefit because, among other reasons, aircraft must be routed for maintenance, and the

Copyright © 1989, The Institute of Management Sciences
0091-2102/89/1904/0020$01.25
This paper was refereed.

TRANSPORTATION — AIR
PROGRAMMING, LINEAR — APPLICATIONS

number of available aircraft is limited. Objectives that can be maximized include utilization of the most efficient aircraft types, operating cost saved, or profits. Although the preferred objective function is to maximize profits, reduced operating cost can be the objective, especially when passenger levels are low enough that traffic is not affected by aircraft capacity.

Operational constraints include the requirement that certain flights operate with specified aircraft types, limits on the number of aircraft that remain overnight at particular stations, and limits on the arrivals or departures (slots) at a station during the day.

The model uses integer linear programming to solve the fleet assignment problem. Given a schedule (with departure and arrival times indicated), it determines which flights should be assigned to which aircraft types to optimize the objective function.

The model can handle both the case where all flights are to be served and the case where some may be dropped. The impetus for the latter case would arise in the early stages of planning a schedule when the schedule contains a "wish list" of extra flights which needs to be pruned to fit the available fleet.

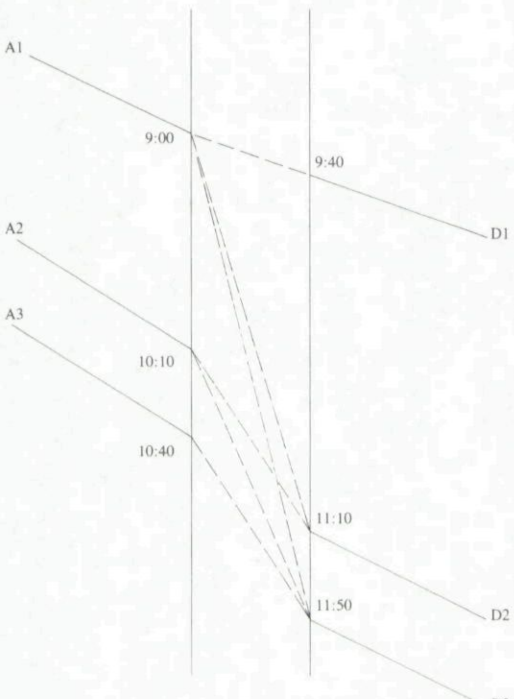

Figure 1: An illustration of feasible turns where A1, A2, and A3 are arriving flights and D1, D2, and D3 are departing flights. Allowing a minimum connection (ground) time of 40 minutes, 12 turn variables per aircraft type are possible: A1-D1, A1-D2, A1-D3, A2-D2, A2-D3, A3-D3, A1-0, A2-0, A3-0, 0-D1, 0-D2, 0-D3. A1-D1, and so forth are arrival-departure turns and A1-0 and 0-D1 are terminating and originating flights.

Formulation

Of the five main groups of constraints, four are intrinsic to the model while the fifth is optional and includes all user-specified rules. The four intrinsic constraints are flight coverage, continuity of equipment, schedule balance, and aircraft count.

Flight Coverage

Each arriving flight may connect with any departing flight whose departure time permits a minimum time (40 minutes) for the connection, unless such a connection is prohibited. Flight-to-flight connections usually are referred to as turns. Typically, an arriving flight can turn to more than one departing flight (Figure 1).

To prevent flights being counted twice, each flight must be limited to being served no more than once. In other words, no more than one of a flight's possible turns can be active. If all flights must be served, the constraint would specify that each flight must be served exactly once.

Continuity of Equipment

It is necessary that each flight served begin (sequence origination or continued from another flight) and end (sequence termination or turn into another flight) on the same aircraft type. This assures the integrity of the network.

Schedule Balance by Station and Aircraft Type

Provision is made for a schedule that is not balanced. An excess of arrivals over departures at a station results in a sequence origination shortage; the reverse

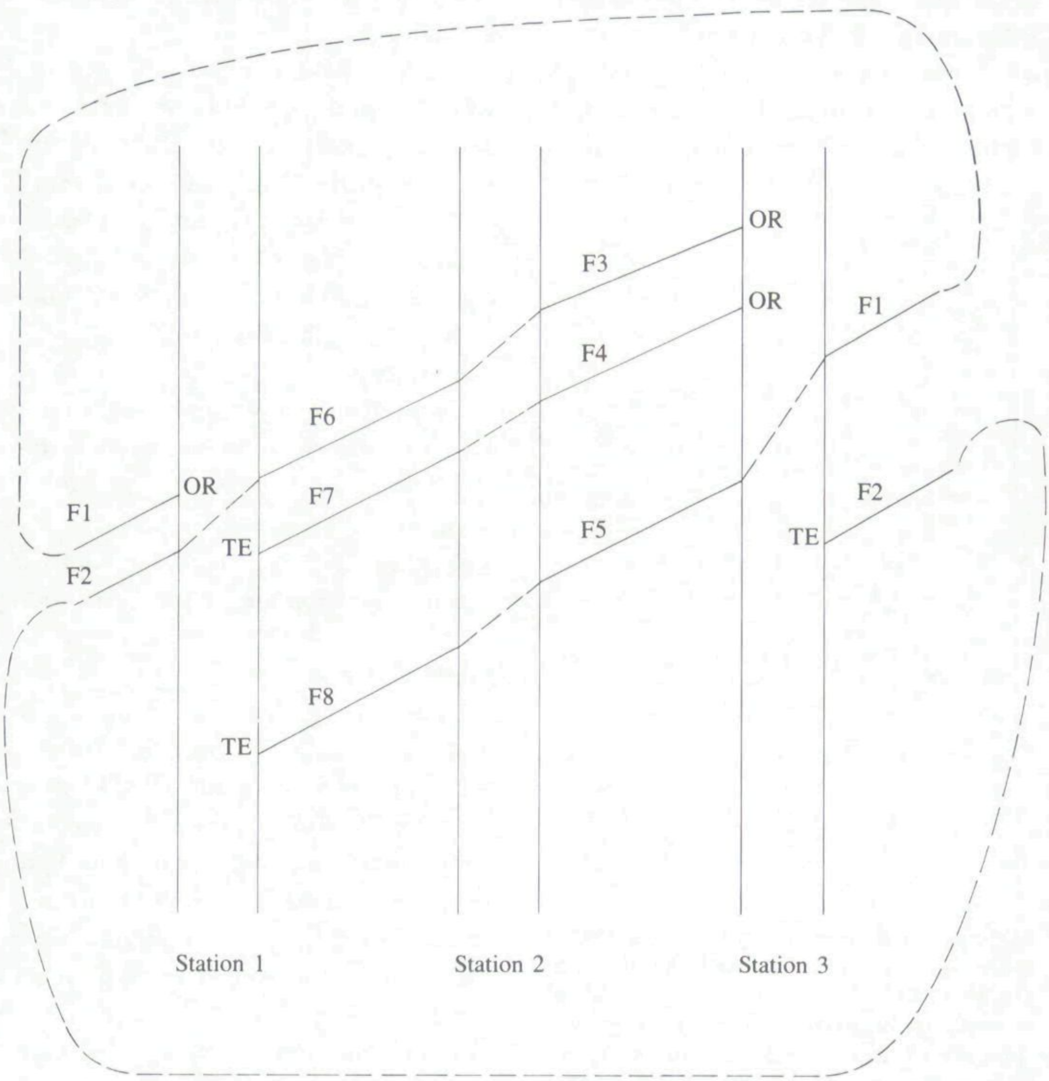

Figure 2: An illustration of sequences with an unbalanced schedule: (1) F1 — F5 — F8, (2) F3 — F6 — F2, (3) F4 — F7. Station 2 is balanced; stations 1 and 3 are not. There is a sequence origination shortage at station 1 and a sequence termination shortage at station 3. (F1 represents flight 1, and so forth; OR is a sequence origination; and TE is sequence termination.)

situation leads to a sequence termination shortage. A difference between the schedule's total departures and total arrivals represents a physical imbalance. But a schedule may be balanced physically and still have imbalances within two or more aircraft types. The balance problem is handled by introducing an origination shortage variable and a termination shortage variable for each station and aircraft type combination. The sum of the sequence originations and the origination shortage variable must equal the sum of the sequence terminations and the termination shortage variable. This is a form of Kirchhoff's Law of Conservation of Flows (Figure 2).

Aircraft Count

A count is kept of the aircraft. A secondary goal of the model is to minimize the number of aircraft used. Therefore, if the schedule is too small for the available aircraft, only the number needed should be used. If the schedule is too large and all flights are to be served, the available aircraft of all types should be exhausted before any extra aircraft are added.

American schedules over 2,300 flights per day to over 150 different cities using 500 jet aircraft.

Fifth Group of Constraints

It is possible to output lower and upper limits on almost any other flight related variable. The only requirement is that the variable must be additive, that is, the value associated with a combination of flights must be equal to the sum of the values of the individual flights. Examples include number of segments, aircraft utilization, and operating cost.

For instance, it may be necessary to place limits on the utilization of the aircraft and on the system cost by aircraft type. Limits also may be needed on the number of aircraft overnighting at an airport or on other additive attributes in the system.

There could be a limit on the daily flights at an airport or on hour movements at the station. In some cases, the flights in a city-pair or market are limited, or there is some limit on the number of flights which may be flown or assigned to particular aircraft types.

The limits may be upper or lower bounds or equalities. Possible constraints include

— Limits on aircraft overnighting at a given station (an overnighting aircraft is equivalent to a sequence termination in a one-day cycle);
— Limits on overnighting aircraft for a group of stations;
— Limits on slots or daily service;
— Limits on system operating costs; and
— Forced turns requiring that a specific in-flight turn into a specified outflight.

There also may be limits on the number of stations served. Sometimes, say for a new aircraft in the system, it may be necessary to limit how many stations are served by a fleet. This may be induced by the fact that overnighting a new aircraft type at a station incurs incremental maintenance personnel and parts inventory costs. There also may be incremental costs associated with some fleet assignments because the locations of the crew bases for that particular fleet are incompatible with the assignments. These station constraints tend to increase computer run times substantially, sometimes making a solution impossible.

Objective Function

Each flight's contribution to the objective function is the value associated with its benefit of interest — profit, aircraft utilization, and so forth. This contribution is assigned arbitrarily to the turns in which the flight is the departing segment.

Aircraft use requires aircraft ownership and that involves physical aircraft with costs that include lease, insurance, or other ownership costs. Shortages in sequence originations and terminations result in dead-heading and incur costs.

The objective is to maximize the benefit contributions of the flights less the cost of aircraft used and the cost of aircraft shortages (imbalances) and the cost of stations. The net effect of this formulation is to maximize the benefit of interest using the smallest number of aircraft possible and minimizing the level of schedule imbalance.

Problem Size and Computation Time

Each feasible turn and aircraft combination represents a decision variable. Each flight segment generates a sequence origination turn, a sequence termination turn, and one or more flight-to-flight turns. The total number of potential flight-to-flight turns at a station is approximately equal to $0.5n^2$, where n is the number of arrivals or departures at a station. Obviously, the average number of turns per flight is higher at the hub stations than at nonhub stations. But for estimating purposes, a typical mix of hub and nonhub stations could yield roughly three flight-to-flight turns per flight. Therefore, each flight would generate approximately five turn variables. Assuming that each flight is permitted on all aircraft types results in a possible total of $5FK$ variables, where F and K are the number of flights and aircraft types, respectively.

There also are the balance or shortage variables, two at each station for each aircraft type for a total of $2KS$ where S is the number of stations. Finally, there are K extra aircraft variables.

The basic rows would generally comprise F flight coverage equations, FK continuity of equipment constraints, KS balance equations, and K aircraft count constraints. User-specified rules are additional constraints.

The problem size can get very large even for medium-sized schedules. For example, a 400-flight schedule with 60 stations and three aircraft types would involve approximately 6,300 columns and 1,800 rows in the ensuing LP matrix. The size of the problem is compounded by the fact that the decision variables must be integer, requiring that an integer linear programming algorithm be invoked.

In practice, two things have been noted. First, because of the routing cycles inherent in the assignments, it is not necessary to explicitly require that the balance variables be integer during the ILP phase. It is sufficient to require only that the turn variables be integer. Of course, the latter are the bulk of the decision variables. Second, the continuous solution usually is integer or is fractional for only a few flights so that an integer solution can be found in a few steps. However, finding an integer solution sometimes can require many ILP iterations and result in lengthy computation times.

Using the IBM MPSX/370 and MIP/370 on an IBM 3081 machine, run times have ranged from slightly under two minutes for a two-aircraft-type problem to over 60 minutes for a problem involving four aircraft types. Most of the runs for two-aircraft and three-aircraft type problems have been in the 15 to 30 minute range. Run times tend to increase faster than linearly with increases in the total number of the LP matrix elements (columns plus rows).

Inputs

The prohibition of turns alluded to earlier can be achieved either by explicitly eliminating a turn or by imposing a penalty on it in the objective function. In the latter case, the flight turn's contribution is modified as follows:

$$\text{Revised contribution} = \text{Original contribution} - \text{Penalty.}$$

To promote user-friendliness, we made it possible to eliminate a turn or impose a penalty by indicating flight number, city pair, station, stage length, or historical load, each within an aircraft type. Using penalties instead of eliminating flight variables allows for estimating the costs of different decisions, such as restricting service in a market to a particular aircraft type.

Applications

American Airlines has used the model thus far for ad hoc studies for one-time decisions. The studies have covered a broad range of areas including fleet planning, crew base planning, and schedule development.

Crew Base Analysis

The model is used as one of two key modules imbedded in American's crew base planning system. The system allows management to analyze the cost of a wide variety of crew-base scenarios where the decision variables are (1) Where should American have a crew base? (2) What type of aircraft should American operate out of that crew base? and (3) How large should the crew base be? Without the fleet assignment model the crew-base planning system could not have been developed and only limited analysis evaluating only a fraction of the alternatives would be feasible.

Cost Reduction

The model optimizes one primary objective at a time, but the constraints and individual flight contributions to the primary objective can be adjusted so that it addresses secondary objectives along with the primary one. In this application, we used the model to minimize the operating costs for the given set of flights. At the same time, we used biases to tempt the model to assign the larger aircraft to the high load flights. We used historical data to determine which flights were high load.

On the basis of traffic, we rated the top 25 percent of the legs as high and set a very high cost penalty to prohibit the small aircraft from being assigned to those legs unless absolutely necessary. We rated the bottom 10 percent of the legs as low and imposed a smaller cost penalty if they were served with the larger aircraft. This discouraged such assignments unless they were absolutely necessary to achieve one of the appropriate high leg assignments.

The number of high legs that were covered by the larger aircraft increased from 76 percent to 90 percent compared to the routings initially proposed. At the same time, operating cost was reduced 0.5 percent. After-the-fact review suggested that the revenue gained from the additional high legs being covered by larger aircraft was equivalent to approximately one percent of revenue. American's 1988 revenues were in excess of $7.5 billion.

Profit Improvement

We have used the model in an exercise to maximize operating profit based on traffic estimates. Operating profit is defined as the difference between expected revenue and operating costs.

We imposed operational constraints, including limits on overnighting aircraft and prohibition of specified aircraft types from certain stations and markets. All flights were served, and four aircraft types were involved.

Using the initial assignments of aircraft to the same schedule as a basis, we reduced operating costs by 0.4 percent and increased the operating margin by 1.4 percent. We also decreased the number of flight segments in which the aircraft with the highest profit was not selected because of routability by 15 percent.

Aircraft Utilization Maximization

We used the model to reduce the cost of flying the schedule by maximizing the utilization of our most cost-efficient narrow-body fleet (MD80) at the expense of our older, less efficient 727 fleet. The model was able to increase the average daily utilization of the MD80 fleet by more than one hour per day.

Conclusions

The model described in this paper has evolved over the past six years to become a very useful tool for one-time decision support projects and has affected decision making at American Airlines in a very significant way. However, the future of the fleet assignment model is even more important because it is a key decision module being incorporated into the next generation scheduling system currently being developed at American Airlines. As part of the new system, the model's role will be increased from an ad hoc decision making tool to a tool used daily by schedule analysts to develop American's future and current schedules.

APPENDIX

Let

X_{ijk} = Feasible turn (flight leg i turns to flight leg j on aircraft type k; if $i = 0$, then j is a sequence origination; if $j = 0$, then i is a sequence termination; where a sequence represents the daily routing for an aircraft),

e_k = Extra aircraft of type k used beyond number specified,

M_k = Available aircraft of type k,

P_{jk} = Benefit (or profit) of operating flight j on aircraft type k,

O_{sk} = Sequence origination shortage of aircraft type k at station s,

T_{sk} = Sequence termination shortage of aircraft type k at station s,

Y_{sk} = Indicator of service/no service of aircraft type k at station s,

C_1 = Nominal cost per aircraft used (typical value = 1),

C_2 = Large cost per extra aircraft (typical value = 800,000),

C_3 = Large cost per imbalance (shortage) (typical value = 500,000),
F = Number of flights,
K = Number of aircraft types,
S = Number of stations,
A_s = Set of arrivals at station s,
D_s = Set of departures from station s,
AD_s = Combined set of arrivals at and departures from station s, and
CS_k = Imposed cost (penalty or reward) for each station served by aircraft type k,

with

$X_{ijk} = 0, 1,$
$Y_{sk} = 0, 1,$
$O_{sk}, T_{sk} = 0, 1, 2, \ldots .$

Constraints
Flight Coverage: This constraint states that every flight served must be a sequence origination or continued from another flight and for only one aircraft type.

$$\sum_{i=0}^{F} \sum_{k=1}^{K} X_{ijk} \leq 1 \text{ for all } j. \tag{1}$$

Continuity of Equipment:

$$\sum_{i=0}^{F} X_{ilk} = \sum_{j=0}^{F} X_{ljk} \quad \text{for all } l, k. \tag{2}$$

Schedule Balance:

$$\sum_{i \in D_s} X_{0ik} + O_{sk} = \sum_{i \in A_s} X_{i0k} + T_{sk} \quad \text{for all } s, k. \tag{3}$$

Aircraft Count:

$$\sum_{i=1}^{F} X_{0ik} - e_k = M_k \quad \text{for all } k. \tag{4}$$

Some Operational Constraints:

Limits on overnighting aircraft for a group of stations:

$$100 \sum_{s \in G_k} \sum_{i \in A_s} X_{i0k} \geq L_k \sum_{i \in A_s} X_{i0k}, \tag{5}$$

or

$$(100 - L_k) \sum_{s \in G_k} \sum_{i \in A_s} X_{i0k} - L_k \sum_{s \notin G_k} \sum_{i \in A_s} X_{i0k} \geq 0 \tag{6}$$

for each
affected k,

where G_k is the group of stations for fleet k, and L_k is a lower bound on the percent of overnights in the fleet which must come from the group of stations; an overnight is equivalent to a sequence termination.

Slot limits for stations by time of day:

$$\sum_{k=1}^{K} \sum_{i \in A_s} \sum_{j=0}^{F} X_{ijk} / (t_1 \leq A_i \leq t_2) \leq U \tag{7}$$

for each affected k
and each time
interval $t_1 - t_2$

where the limit is for arrivals; A_i is the arrival time, t_1 and t_2 define the time interval, U is an upper bound and the limit is for all aircraft types.

A limit on the number of stations served can be achieved by either or both of two ways: specifying a numerical limit on the number of stations or imposing a cost on each station served. These station constraints tend to increase run times substantially, sometimes making a solution impossible.

The corresponding constraints are as follows:

(1) The number of flights *actually flown* (or assigned) into and out of a station cannot exceed the number *permitted* at the station, by aircraft type.

$$\sum_{i,j \in AD_s} X_{ijk} \text{(assigned)} \leq Y_{sk} \sum_{i,j \in AD_s} X_{ijk} \text{(permitted)} \tag{8}$$

if $LS_k > 0$
or $CS_k \neq 0$,

where Y (= 0,1) indicates no service/service of aircraft type k at station s; LS_k is the bound (upper or lower) on the number of stations to be served by aircraft type k; and CS_k is an imposed cost (penalty or reward) for each station served by aircraft type k. A positive CS represents a penalty which minimizes the number of stations, while a negative CS will tend to increase the number of stations served.

(2) If a station is served by aircraft type k, then there must be at least one flight from or into it within the aircraft type.

$$\sum_{i,j \in AD_s} X_{ijk} \geq Y_{sk} \quad \text{if } LS_k > 0 \text{ and represents an exact value or a lower bound or } CS_k \neq 0. \tag{9}$$

(3) This constraint applies only if the limit on stations is effected by an explicit bound rather than by an imposed cost.

$$\sum_s Y_{sk} \begin{Bmatrix} \leq \\ = \\ \geq \end{Bmatrix} LS_k \text{ if } LS_k > 0. \tag{10}$$

Objective Function:
Maximize

$$\begin{aligned} Z = & \sum_{i=0}^{F} \sum_{j=0}^{F} \sum_{k=1}^{K} P_{jk} X_{ijk} \\ & - C_1 \sum_{i=1}^{F} \sum_{k=1}^{K} X_{0ik} \\ & - C_2 \sum_{k=1}^{K} e_k \\ & - C_3 \sum_{s=1}^{S} \sum_{k=1}^{K} (0_{sk} + T_{sk}) \\ & - \sum_k CS_k \sum_s Y_{sk}. \end{aligned} \tag{11}$$

Thomas M. Cook, American Airlines Decision Technologies, PO Box 619616, Dallas/Fort Worth Airport, Texas 75261-9616, writes "The purpose of this letter is to certify that the fleet assignment model discussed in Jeph Abara's paper has been used successfully at American for a number of ad hoc studies and has influenced important decisions at American Airlines. In addition, the fleet assignment model will be a key decision module of the 'Next Generation Scheduling System' currently being developed here at American."

Copyright 1989, 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.