A Lagrangian Relaxation Approach to Assigning Aircraft to Routes in Hub and Spoke Networks

MARK S. DASKIN and NICHOLAOS D. PANAYOTOPOULOS

Northwestern University, Evanston, Illinois 60208

The problem of assigning aircraft to routes to maximize profits in a hub and spoke network is formulated as an integer linear programming problem. A Lagrangian relaxation of the problem is outlined together with heuristics for converting the Lagrangian solutions into primal feasible solutions and for improving on the solutions. Computational results from 324 runs spanning a range of problem sizes are reported. The results suggest that the Lagrangian relaxation is effective at providing an upper bound on the profits and the heuristics yield good solutions when the maximum number of aircraft required to fly all routes in the schedule is less than or equal to the number of available aircraft.

INTRODUCTION

A critical determinant of an airline's survival in today's competitive environment is its ability to utilize its resources efficiently. Critical resources available to an airline include its personnel or crews, its gates and airport slots and its aircraft. In this paper, we introduce a model that can assist airline planners with hub and spoke networks in deploying their fleets as efficiently as possible. Specifically, we outline an optimization model that assigns a fleet of aircraft of different types to routes to maximize profits.

Following deregulation in the United States, the route structure of many domestic carriers evolved from trunk networks to hub and spoke systems.[12] Many European carriers such as Sabena of Belgium, Swissair, and Olympic Airways of Greece employ hub and spoke networks for their European flights. The model outlined below is directly applicable to airlines with a single hub from which all routes depart and to which all routes return. The extension of this work to more general route structures is an important area of research.[4, 9, 15] However, as discussed below, the model is applicable in a large number of other facility utilization contexts including classroom assignment problems.

In many hub and spoke airline planning contexts, the assignment of aircraft to routes follows the determination of: (1) which cities to serve, (2) the frequency of service between the hub and each of the other cities, (3) the desired departure and arrival times into and out of the hub for each route by aircraft type, and (4) the potential profit associated with assigning each route to each different type of aircraft in the fleet. For international carriers, many of these decisions are heavily constrained by bilateral agreements, the availability of gate space in foreign airports, airport curfews and operating regulations, and so on.[11] Thus, in practice, there are relatively few options available to planners with regard to these decisions. The primary source of leverage for such carriers in the short term is their ability to choose which aircraft to assign to each route.

The remainder of this paper is organized as follows. In Section 1, we formulate the aircraft assignment problem as an integer programming model. In Section 2, we outline a Lagrangian procedure for solving the problem. The Lagrangian procedure turns out to be useful in providing an upper bound on the objective function, but not in determining a feasible solution to the original unrelaxed problem. Thus, a heuristic procedure for converting the infeasible Lagrangian solution into a feasible assignment of aircraft to routes and then improving on the assignment is also outlined in Section 2. Section 3 describes an experimental design used to test the computer implementation of the algorithms and presents the results. Directions for future study are outlined in Section 4.

1. PROBLEM FORMULATION

THE PROBLEM with which we are concerned is that of assigning aircraft to routes to maximize the profit to

Transportation Science
Vol. 23, No. 2, May 1989

0041-1655/89/2302-0091 $01.25

be obtained from the scheduled routes. For our purposes, we assume that each route originates at a single hub, visits a number of other cities (usually only one) and returns to the hub. Associated with each route is a scheduled departure and arrival time at the hub as well as the expected profit that will result from assigning the route to each aircraft in the fleet. We also have the required ground servicing time for each candidate route.

We define the following indices and inputs:

i = index of routes
I = set of routes
j = index of aircraft
J = set of aircraft
k = index of time periods
K = set of time periods
D_i = departure time of route i from the hub
A_i = arrival time of route i at the hub
GT_i = ground time required for processing route i
A_i' = effective arrival time of route i at the hub (i.e., the time at which an aircraft assigned to route i would be available for assignment to another route)
$= A_i + \mathrm{GT}_i$
P_{ij} = profit due to assigning route i to aircraft j
$\mathbf{N}_{jk}$ = set of routes i that would utilize aircraft j during period k if assigned to aircraft j.

In addition, we define the following decision variable:

$$X_{ij} = \begin{cases} 1 & \text{if route } i \text{ is assigned to aircraft } j \\ 0 & \text{if not.} \end{cases}$$

The problem of assigning aircraft to routes to maximize profits may be formulated as follows[11]:

$$\text{Maximize} \quad \textstyle\sum_i \sum_j P_{ij} X_{ij} \qquad (1)$$

$$\text{Subject to:} \quad \textstyle\sum_j X_{ij} \leq 1 \quad \forall i \qquad (2)$$

$$\textstyle\sum_{i \in \mathbf{N}_{jk}} X_{ij} \leq 1 \quad \forall j, k \qquad (3)$$

$$X_{ij} = 0, 1 \quad \forall i, j. \qquad (4)$$

The objective function (1) maximizes the profit associated with the assignment. Constraint (2) ensures that each route i is assigned to at most one aircraft. Constraint (3) guarantees that each aircraft j is assigned to at most one route during each time period k. In general, a new time period should be defined whenever there is a chance of needing an additional aircraft.

We employ an inequality for constraint (2) to ensure that a feasible solution exists in the event that the number of available aircraft is less than the number needed to assign all routes to an aircraft. It is also possible to construct examples in which profits are reduced by requiring all routes to be flown. Table I illustrates such a case for a problem with 4 routes and 3 aircraft. If all routes must be flown, routes 1 and 4 must be flown by the same aircraft since all other pairs of routes overlap in time. The combined profit from these two routes is 40, independent of which aircraft is assigned to the routes. Assigning aircraft 2 to route 2 maximizes the profit from route 2. Thus, one profit maximizing assignment with all routes flown is to assign routes 1 and 4 to aircraft 1; route 2 to aircraft 2; and route 3 to aircraft 3 for a total profit of 75. If we allow route 3 not to be flown, the profit maximizing assignment is to assign route 1 to aircraft 1; route 2 to aircraft 2; and route 4 to aircraft 3 for a total profit of 90 units.

A number of alternative approaches may be used to define time periods.[3] The approach that results in the smallest number of time periods, and hence the smallest number of constraints, is to initiate a new time period with each departure that follows an arrival. Arrivals need not initiate a new time period. Constraints associated with a time interval between an arrival and a subsequent departure will be redundant, as all aircraft that are busy during such a time interval must also have been busy during the time period immediately preceding the arrival. Similarly, a departure that immediately follows another departure need not initiate a new time period, since the constraints associated with the first time interval will be redundant with those of the subsequent time interval. Thus, time periods which begin due to a departure which immediately follows another departure can be merged into the earlier time period. The maximum number of time periods generated by this approach will be $|\mathrm{I}|$, though considerably fewer may be expected. In the example shown in Table I, only two time periods are generated using this approach. Table II summarizes the resulting time periods and illustrates constraints (3) for aircraft 1.

TABLE I
Example Problem

Route	Departure Time	Arrival Time	Profit		
			A/C 1	A/C 2	A/C 3
1	0800	1200	30	20	10
2	1000	1400	20	30	10
3	1100	1800	5	5	5
4	1500	2100	10	20	30

TABLE II
Time Periods and Constraints (3) for Aircraft 1 in Example Problem

Period	Begins	Ends	Constraints
1	0800	1500	$X_{11} + X_{21} + X_{31} \leq 1$
2	1500	2100	$X_{31} + X_{41} \leq 1$

2. SOLUTION ALGORITHM

THE PROBLEM formulated in Section 1 is a large integer programming problem. As noted above, we could have up to $|I|$ time periods. This would result in a problem with $|I| \cdot |J|$ decision variables and up to $|I|(|J| + 1)$ constraints. For a problem with 200 routes and 25 aircraft, we would have 5000 integer variables and up to 5200 constraints. In the remainder of this Section, we outline a Lagrangian based procedure for obtaining an upper bound and a heuristic for obtaining good feasible solutions to the problem defined in Section 1.

2.1. A Lagrangian Bound on the Problem

We relax constraint (3) to obtain the following Lagrangian problem:

Optimize $L(\underset{\sim}{X}, \underset{\sim}{\beta})$

$$= \sum_i \sum_j P_{ij} X_{ij} + \sum_j \sum_k \beta_{jk} (1 - \sum_{i \in \mathbf{N}_{jk}} X_{ij}) \quad (4a)$$

Subject to: $\sum_j X_{ij} \leq 1 \quad \forall i \quad (2)$

where

β_{jk} = Lagrange multiplier associated with constraint (j, k) of constraint set (3).

The objective function (4a) can be rearranged as follows:

Optimize $L(\underset{\sim}{X}, \underset{\sim}{\beta})$

$$\begin{aligned} &= \sum_i \sum_j (\mathrm{P}_{ij} - \sum_{k \in \mathbf{M}_{ij}} \beta_{jk}) \mathrm{X}_{ij} + \sum_j \sum_k \beta_{jk} \\ &= \sum_i \sum_j \hat{\mathrm{P}}_{ij} \mathrm{X}_{ij} + \sum_j \sum_k \beta_{jk} \end{aligned} \quad (4b)$$

where

$\mathbf{M}_{ij}$ = set of time periods k during which aircraft j would be busy if it were assigned to fly route i

$\hat{P}_{ij}$ = the "pseudo-profit" associated with route i being assigned to aircraft j
$= P_{ij} - \sum_{k \in \mathbf{M}_{ij}} \beta_{jk}$.

For fixed values of the Lagrange multipliers, β_{jk}, the Lagrangian problem decomposes into subproblems for each route i. Each subproblem may be written as follows:

Sub-Problem i

Maximize

$$\sum_j (P_{ij} - \sum_{k \in \mathbf{M}_{ij}} \beta_{jk}) X_{ij} = \sum_j \hat{P}_{ij} X_{ij} \quad (5)$$

Subject to: $\sum_j X_{ij} \leq 1. \quad (6)$

Each of these subproblems may be solved by inspection. We simply find the maximum pseudo-profit from among all pseudo-profits associated with route i. If this value is positive, we set the corresponding value of $X_{ij} = 1$ and all other values of $X_{ij} = 0$ (for the given route i). If the pseudo-profits for two or more aircraft and a given route are equal and positive, the route is assigned to the aircraft which would result in violating as few of the relaxed constraints (3) as possible. If the maximum pseudo-profit is nonpositive, we set all values of $X_{ij} = 0$ (for the given route i). Note that since the solution to the subproblems is guaranteed to be integer, the Lagrangian upper bound will not be any tighter than would the bound obtained from the linear programming relaxation.[5, 6]

If the values of X_{ij} computed from the relaxed problem at iteration n do not violate any of the relaxed constraints, we can compute a lower bound on the original objective, LB^n, as follows:

$$\mathrm{LB}^n = \max(\mathrm{LB}^{n-1}, \sum_i \sum_j P_{ij} X_{ij}^n) \quad (7)$$

where subscripts are used to indicate the iteration number. The initial lower bound, LB^0, is set to 0. If the values of X_{ij} violate any of the relaxed constraints, we set $\mathrm{LB}^n = \mathrm{LB}^{n-1}$.

The Lagrangian objective function (4b) provides an upper bound on the objective function (4) for the original problem. Since the value of (4b) is not monotonically decreasing, we set the upper bound on iteration n of the algorithm, UB^n, to:

$$\mathrm{UB}^n = \min [\mathrm{UB}^{n-1}, L(\underset{\sim}{X}^n, \underset{\sim}{\beta})]. \quad (8)$$

The Lagrange multipliers, β_{jk}, are updated using a subgradient optimization procedure[4, 5] along with the damping option suggested by CROWDER.[2] Specifically, for each Lagrange multiplier, a search direction, d_{jk}^n, is computed on the nth iteration as follows:

$$d_{jk}^n = (1 - \sum_{i \in \mathbf{N}_{jk}} X_{ij}) + \gamma d_{jk}^{n-1} \quad (9)$$

where

γ = Crowder's damping factor $(0 \leq \gamma < 1)$.

The step size, t^n, at each iteration is computed using the following relation:

$$t^n = \frac{\theta^n (\mathrm{UB}^n - \mathrm{BS}^n)}{\sum_j \sum_k (d_{jk}^n)^2} \quad (10)$$

where

BS^n = best solution obtained so far (i.e., the maximum of the lower bound and the best heuristic solution obtained so far using the procedures outlined below)

θ^n = a multiplier whose value is halved when the number of successive iterations during which we fail to improve the upper bound, UB, exceeds some specified value. The initial value of θ is set to 2.

Finally, the Lagrange multipliers for the next iteration are computed using the following relation:

$$\beta_{jk}^{n+1} = \max\,[0,\ \beta_{jk}^{n} - t^{n} d_{jk}^{n}]. \qquad (11)$$

2.2. A Heuristic Algorithm

As noted above, if the Lagrangian solution satisfies all of the relaxed constraints (3), the solution is feasible for the original problem. In this case, we can use the value of the objective function (1) to update the lower bound as indicated in Equation 7. In our experience, this almost never happens. This is not surprising in light of the very large number of relaxed constraints all of which must be satisfied for the Lagrangian solution to be primal feasible.

To deal with the likely infeasibility of the solution to the relaxed problem, we developed a heuristic procedure to convert an infeasible Lagrangian solution into a feasible primal solution. Other heuristics were developed to improve upon this solution. Figure 1 is a macro-flowchart of the entire solution procedure, including the Lagrangian algorithm. In the remainder of this section, we focus on the heuristic procedures used to identify feasible solutions based on the relaxed Lagrangian solutions.

Before discussing the heuristics, we assume that the aircraft may be grouped into a small number of aircraft types, based on the profit associated with the route and the aircraft. The profit associated with assigning a route to an aircraft will be the same for all aircraft of the same type. In what follows, we let

t = index of aircraft types
T = set of aircraft types
P_i^t = profit due to assigning route i to an aircraft of type t.
NP_i = number of periods associated with route i.

We note that relaxing this assumption is relatively straightforward.

If the Lagrangian solution is not feasible, it is because some aircraft are assigned to multiple routes at the same time. This infeasibility may be resolved either (a) by *reassigning* all but one of the routes that are assigned to such aircraft during a time period to other aircraft that are available to serve these routes or (b) by "*unassigning*" some of the conflicting routes assigned to the same aircraft in the Lagrangian solution. Recall that constraint (2) allows routes not to be assigned to any aircraft. Specifically, we examine each aircraft in turn. For each aircraft j, we consider each time period, k, during the planning period. If the aircraft is assigned to multiple routes during period k, we find the route from the set of routes assigned to the aircraft during period k which minimizes the difference between the profit of assigning the route to aircraft j and assigning it to another aircraft, j', which is available to fly that route. If such a reassignment can be identified we make this reassignment, thereby reducing by 1 the number of conflicting routes to which aircraft j is assigned during period k.

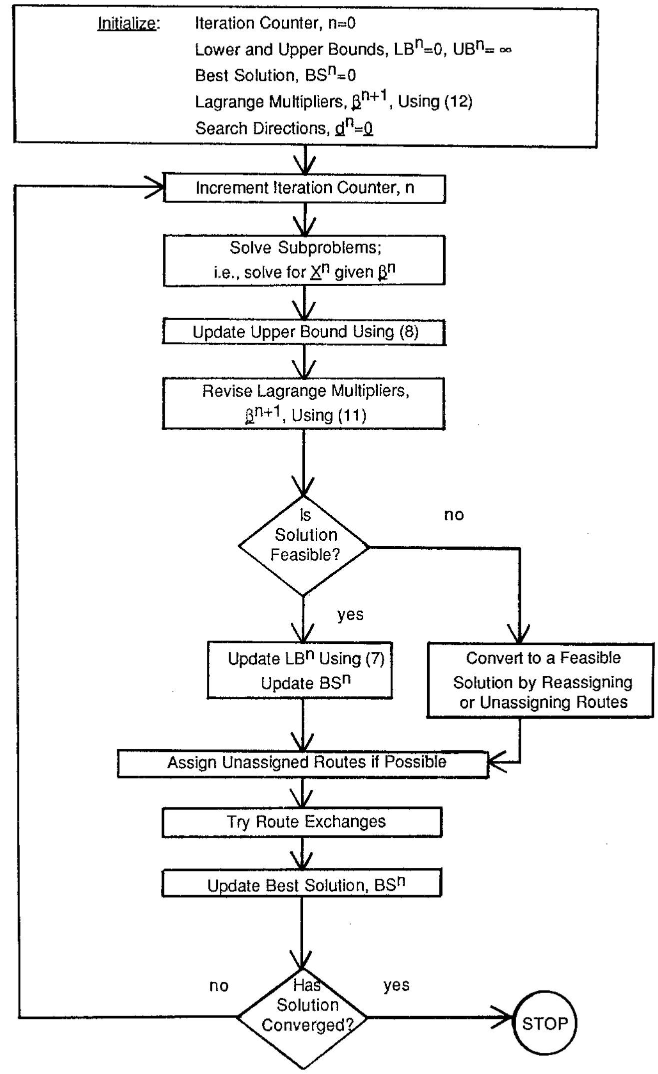

Fig. 1. Flowchart of the solution procedure.

If no route can be reassigned, we unassign a route from among the routes currently assigned to aircraft j during period k. This too will reduce by 1 the number of conflicting routes to which aircraft j is assigned during period k. The route to be unassigned is the one which minimizes the following quantity:

$$\mathrm{OP}_i = \begin{cases} \dfrac{\max_t (P_i^t) - \max_{t \neq t_i^*} (P_i^t)}{\mathrm{NP}_i} & |T| \geqslant 2 \\[2ex] \dfrac{P_i^1}{\mathrm{NP}_i} & |T| = 1 \end{cases}$$

where

OP_i = "opportunity cost" associated with route i
$t_i^* = \underset{t}{\operatorname{argmax}}\ (P_i^t)$
= aircraft type which maximizes the profit for route i.

Picking such a route to be unassigned is an attempt to identify a route for which there are likely to be a large number of good aircraft assignments to consider during later phases of the algorithm. For such routes, the difference between using the best aircraft type and the second best type is small. Normalizing this difference by the number of periods the route requires results in the algorithm tending to unassign long routes. Doing so is likely to reduce the number of conflicts during contiguous time periods, thereby reducing the subsequent need to unassign routes. Clearly, alternative heuristics for selecting routes to unassign can and should be developed and tested.

After all aircraft and all time periods have been examined and conflicting routes have been reassigned to other available aircraft or have been unassigned, the resulting assignment is primal feasible (by construction). Next we try to assign those unassigned routes in the heuristic solution to aircraft. Recall that routes may be unassigned either due to the operation of the heuristic as described above or because all pseudo-profits for the route were negative. In either case, it may be possible to identify routes which can be assigned to available aircraft thereby increasing the total profit associated with the assignment. Processing each unassigned route in turn, we find the available aircraft which maximizes the profit for the unassigned route, if there are any aircraft available to fly the route. If such aircraft can be found, we assign the route to the aircraft which maximizes the contribution to profit.

Finally, we consider aircraft exchanges for all routes that are unassigned or that are not assigned to an aircraft of the type which maximizes the profit for the route. We begin by trying to (re)assign each such route to an aircraft of the type which maximizes the profit for the route. If this can be one, the (re)assignment is made. If it cannot be done, the route is added to a list of suboptimally assigned routes. After considering all such routes, we examine all pairs of suboptimally assigned routes to see if exchanging the aircraft assignments is feasible and profitable. If it is, we do the exchange. In essence, this final phase of the heuristic is like a restricted version of a single pass through a two-opt exchange for a traveling salesman problem.[10]

2.3. Initial Lagrange Multipliers and Stopping Criteria

In all runs reported below, the Lagrange multipliers were initialized to the same value as follows:

$$\beta_{jk}^1 = \max_{i,s}\ (P_{is}) \quad \forall j, k. \tag{12}$$

That is, the initial Lagrange multipliers were set equal to the largest profit for any route/aircraft combination. Since the pseudo-profit for each route/aircraft combination is given by

$$\hat{P}_{ij} = P_{ij} - \textstyle\sum_{k \in \mathbf{M}_{ij}} \beta_{jk}$$

all pseudo-profits on the first iteration will be nonpositive. Thus, no routes will be assigned to aircraft on the first iteration. The initial upper bound will become

$$UB^1 = \textstyle\sum_j \sum_k \beta_{jk}^1 = |J| \cdot |K| \cdot [\max_{i,j}\ (P_{ij})].$$

Since no routes will be assigned to aircraft on the initial iteration, the solution at the end of the first iteration will be that which results purely from the use of the heuristic algorithms. Also, the new search directions will all be equal to 1, and the revised values of the Lagrange multipliers, $\underline{\beta}^2$, are very likely to be zero using Equations 9, 10, and 11. Thus, on the second iteration, the pseudo-profits for each route are likely to equal the route's real profit. Thus, the upper bound on the second iteration will be (in all likelihood)

$$UB^2 = \textstyle\sum_i \max_j\ (P_{ij}), \tag{13}$$

or the profit that would result if each route could be assigned to the most profitable aircraft for that route. Note that we cannot be certain of the values of the Lagrange multipliers on the second iteration as they depend on the step size, t^1, which in turn depends on the relative magnitudes of the upper bound, UB^1, and the heuristic value, BS^1.

Another way of initializing the Lagrange multipliers would be to set them all to zero ($\beta_{jk}^1 = 0\ \forall j, k$). This approach results in initial search directions which are nonpositive (9) and initial increases in all Lagrange multipliers. Our experience suggests that this in turn results in an inability to reduce the upper bound below that given by (13) for many iterations. Thus, the initial Lagrange multipliers were set using Equation 12.

As shown in Figure 1, the algorithm iteratively solves the Lagrangian problem, revises the lower and upper bounds as well as the Lagrange multipliers, and applies the heuristic procedures outlined above. The iterative procedure was allowed to terminate if any one of a number of conditions was satisfied. First, if all routes are assigned to an aircraft of the type which maximizes the route's profit, the assignment is clearly optimal. Second, the algorithm was terminated if θ^n fell below a prespecified value. Third, if the upper bound was less than 1% greater than the value of the best solution found so far, the algorithm stopped. Finally, a limit on the number of iterations was specified.

3. EXPERIMENTAL ANALYSIS

THE ALGORITHM described in Section 2 was coded for an IBM Personal Computer/AT (running at 6 MHz) using Turbo Pascal, version 3.01A. In all runs, θ was halved after 4 successive failures to reduce the upper bound; the minimum allowable value of θ was 5×10^{-7}; Crowder's damping term was set to 0.25; and a maximum of 100 iterations were permitted per case.

To test the algorithm an experimental design involving 324 runs was established. Table III summarizes the key factors in the experimental design and the levels of the factors.

Schedules were generated by randomly selecting a departure time for each route between 0 and 48 hours in increments of 15 minutes. The route duration (i.e., the round trip flight time and any required ground time) was then sampled from one of the two uniform distributions of route durations and truncated to the nearest 15-minute duration. Finally the profit per hour was randomly sampled from a uniform distribution for each aircraft type. The process was repeated until the desired number of routes was obtained for each case. Three independent replications of each combination of factor inputs were performed. In all cases, 24 aircraft were available for assignment to routes.

The number of aircraft needed to allow all routes to be flown may easily be determined for any case by identifying the time and period during which the maximum number of routes are in service. Table IV(a) gives the average number of aircraft needed as a function of the number of routes and the mean route time. As expected, the number of required aircraft increases with both factors. Figure 2 plots the distribution of the number of aircraft required for the 324 test cases. Note that for 56 runs, the number of aircraft needed exceeded the number of available aircraft. In these cases, some routes were not assigned to any aircraft in the best solution identified by the algorithm.

One measure of the quality of the solution is the difference between the upper bound and the value of the best solution found by the program. Table IV(b) summarizes the average difference between the upper bound and the best solution as a percent of the best solution. Specifically, it reports the average over all runs for each cell in the table of the following quantity:

$$\frac{\text{Best upper bound} - \text{best solution value}}{\text{Best solution value}} \times 100.$$

If in a particular run, each route was assigned to an aircraft of the optimal type for that route, the upper bound was set equal to the best solution value in the formula above, yielding a 0% difference for the run.

TABLE III

Summary of the Experimental Design

Factor	No. of Levels	Levels		
Mean route duration	3	6, 8, 10 hours		
Distribution of route duration	2	U(0.5 × Avg, 1.5 × Avg)		
		U(0.75 × Avg, 1.25 × Avg)		
Aircraft types	3	2, 3, 4		
		No. of Aircraft Types	Balanced Composition	Unbalanced Composition
Fleet composition (given number of aircraft types)	2	2	12–12	9–15
		3	8–8–8	6–8–10
		4	6–6–6–6	4–5–7–8
Number of routes	3	50, 75, 100		
Profit per route hour	1	U(10, 30)		
Replications	3	Random route durations, starting times, and profits		
Total cases	324			

TABLE IV

Summary of Model Results

Mean Route Time	Number of Routes: 50	75	100	Average
(a) Average Number of Aircraft Needed to Fly All Routes				
6	11.6	15.9	19.7	15.7
8	13.9	18.9	24.3	19.0
10	16.6	22.4	29.1	22.7
Average	14.0	19.1	24.4	19.2
(b) Average Percent Difference between Upper Bound and Best Solution Value				
6	0.05	0.23	0.55	0.28
8	0.12	0.47	3.53	1.37
10	0.26	1.91	13.54	5.24
Average	0.14	0.87	5.88	2.30
(c) Percent of Routes Which Are Poorly Assigned				
6	0.4	2.1	4.9	2.5
8	1.2	3.9	14.6	6.5
10	2.6	11.6	26.2	13.4
Average	1.4	5.9	15.2	7.5
(d) Average Solution Times (Seconds)				
6	4.8	9.9	54.9	23.2
8	6.5	56.3	985.0	349.2
10	12.1	536.8	1,407.5	652.2
Average	7.4	201.0	815.8	341.5

Note: All times are in seconds on an IBM PC/AT computer operating at 6 MHz with code written in Turbo PASCAL, version 3.01A.

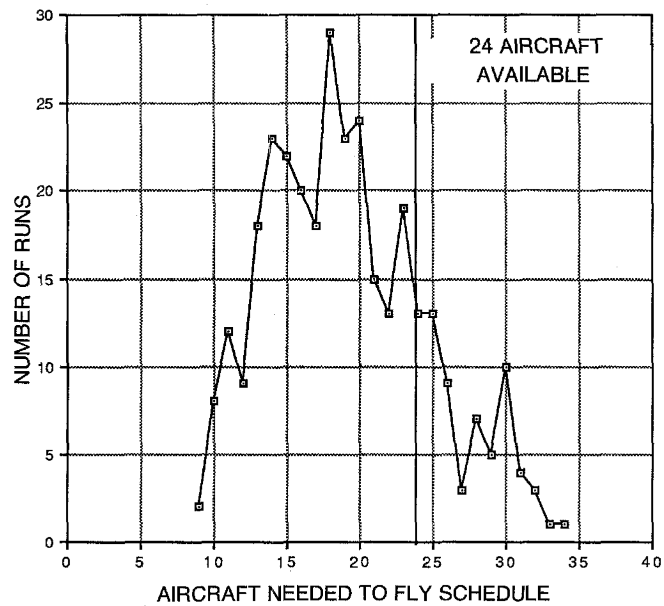

Fig. 2. Distribution of runs versus aircraft needed.

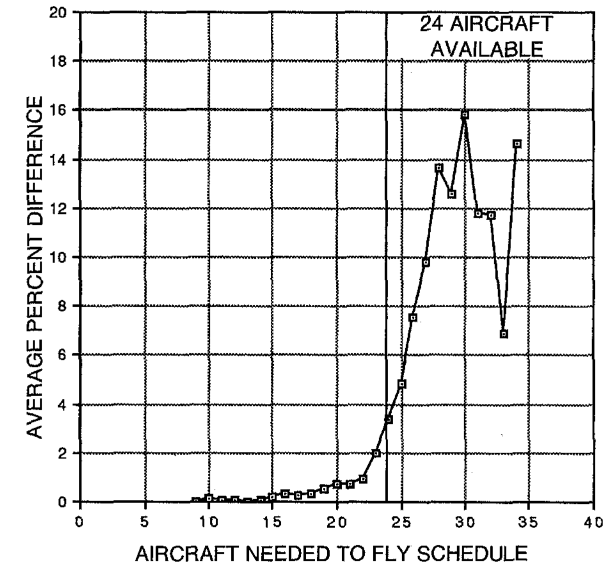

Fig. 3. Average percent difference between upper bound and best solution value.

Figure 3 plots this quantity as a function of the number of aircraft needed to fly all routes. When the number of required aircraft was less than or equal to the number available, the average percentage difference between the upper bound and the best solution value was under 4%. When the number of required aircraft exceeded the available number, the average percentage difference was significantly larger.

Large differences between the upper bound and the best solution value are expected in these cases since the upper bound is no better than the bound obtainable from the linear programming relaxation of (1)–(4), and the LP relaxation of such problems is likely to include a large number of fractional assignments. Table V illustrates this phenomenon for a problem with 3 routes and 1 aircraft. The upper bound would be 18, but the best solution would be 14 for a difference of 28.6%. Thus, additional iterations of the Lagrangian procedure are unlikely to produce a significantly tighter bound in these cases. Embedding the Lagrangian procedure in a branch and bound algorithm may be required to obtain tighter bounds.

A second measure is the percent of routes which are left unassigned or which are not assigned to an aircraft of the type which maximizes the profit for the route. We note that it may be optimal to leave some routes unassigned or to assign others to aircraft of a type that does not maximize the profit for the particular route. Thus, this measure reflects both solution quality (to the extent that an improved solution may exist) and the inherent difficulty associated with solving any particular problem instance. Nevertheless, we refer to such routes as "poorly assigned" routes. Table IV(c) summarizes the average percent of routes which are poorly assigned as a function of the number of routes and the mean route duration. Figure 4 plots this quantity as a function of the number of aircraft required to fly all routes. Again, solution quality goes down as the number of required aircraft increases.

TABLE V

Illustration of the Gap between the LP and IP Solutions

Route	Departure	Arrival	Profit
1	0000	1600	10
2	0800	2400	12
3	1600	0800	14

(24-hour planning period, single aircraft)

LP solution: $X_{11} = X_{21} = X_{31} = 0.5$
Objective Function = 18

IP solution: $X_{11} = X_{21} = 0$
$X_{31} = 1$
Objective Function = 14

Percent difference: $[(18 - 14)/14] \times 100 = 28.6$

Finally, Table IV(d) summarizes the average solution time in seconds as a function of the number of routes and the mean route duration. Figure 5 plots the average solution time versus the number of required aircraft. Solution times are small when there are at least two or three more aircraft available than are required. As the number of aircraft required to fly all routes approaches the number available, solution times increase rapidly. When the number of required aircraft exceeds the available number, solution times are large. Generally these runs terminated because the limit of 100 iterations was reached. The increases in the solution time with the number of required aircraft

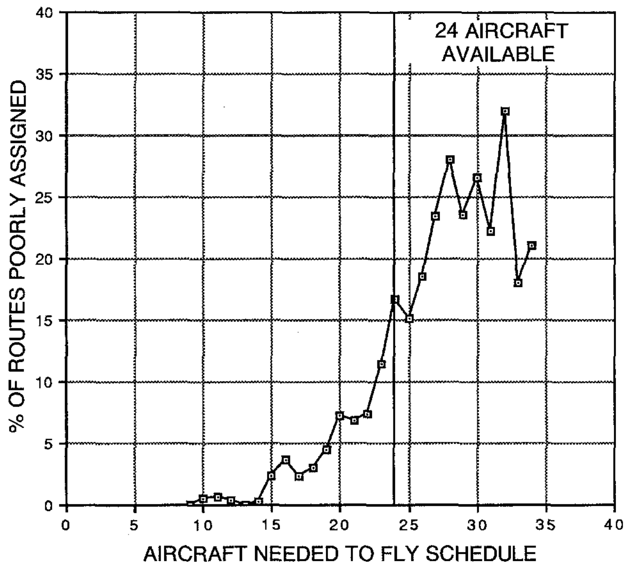

Fig. 4. Percent of routes that are poorly assigned.

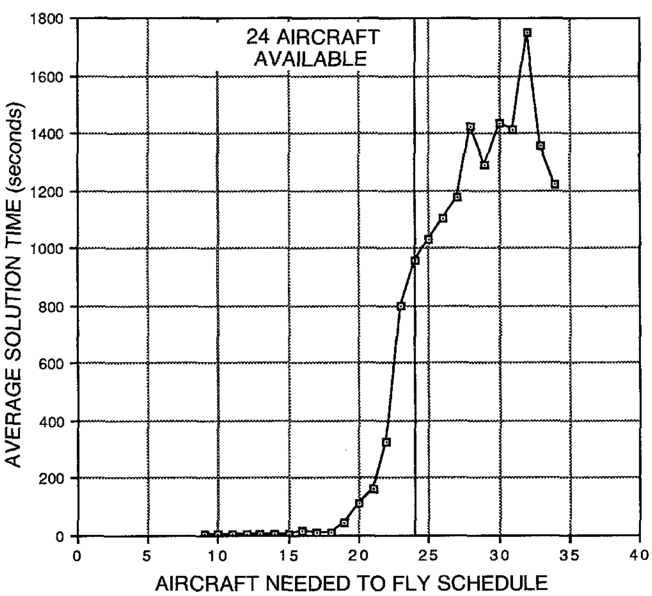

Fig. 5. Average solution time versus aircraft needed.

mirror the degredation in solution quality shown in Figures 3 and 4.

These results suggest that the algorithm is effective at solving the aircraft assignment problem given by Equations 1–4 and that the Lagrangian relaxation provides a good upper bound when the number of aircraft required to fly all routes is less than or equal to the available number of aircraft. When the number of required aircraft exceeds the number available, the difference between the Lagrangian upper bound and the best solution value increases and additional techniques, such as branch and bound, may be required to tighten the bound and/or improve the solution. It is worth noting that when the number of required aircraft is significantly less than the number available, any "greedy" algorithm is likely to yield good results. Conversely, when the number of required aircraft greatly exceeds the number available, we are actually faced with a combined problem of *selecting* routes to be flown and *assigning* aircraft to the selected routes. Not unexpectedly, such cases appear to be quite difficult to solve using the present approach. For intermediate cases in which the number of required aircraft is approximately equal to the number available—the cases most likely to be encountered in practice—the algorithm seems to be effective at solving the aircraft assignment problem.

4. SUMMARY, CONCLUSIONS AND RECOMMENDATIONS FOR FUTURE RESEARCH

WE HAVE formulated an optimization problem to assign aircraft to routes to maximize profits in a single-hub-and-spoke network. A Lagrangian relaxation of the problem was structured to provide an upper bound on the obtainable profits and a variety of heuristics were designed to convert the Lagrangian solution into a primal feasible solution and to improve on the resulting assignment. Tests of the algorithm indicate that it is effective at solving the problem and that the Lagrangian relaxation provides a good upper bound when the number of aircraft required to fly all routes is less than or equal to the number of available aircraft. When the number of available aircraft is less than the number required to fly all routes, the difference between the Lagrangian upper bound and the best assignment value found by the algorithm may be large and the process may need to be embedded in a branch and bound algorithm to improve the bound or the solution value.

The formulation presented in this paper is applicable to a broad range of activity/resource assignment problems. For example, GLASSEY and MIZRACH[7] examine the assignment of classes to rooms at the University of California at Berkeley. Their constraints are similar to (2)–(4) except that their constraint (2) is an equality constraint indicating that all classes must be assigned to a room. Their objective is to minimize the "cost" of the assignment. They present a heuristic solution algorithm only. TRIPATHY[16] discusses a similar classroom scheduling problem and reports the results of using a number of relaxations to the problem. SCHMIDT and STROHLEIN[13] provide a bibliography of papers related to timetable construction problems with an emphasis on the relationship of these problems to graph theory. JUNGINGER[8] surveys the state of the art of computerized classroom scheduling in Germany.

Returning to the application of the model to route/aircraft assignment problems, a number of avenues of future research may be fruitfully pursued. As noted above, the Lagrangian relaxation may be embedded in a branch and bound algorithm. The formulation can easily be extended to allow the required ground time for a route to depend on the aircraft assigned to fly the route.[3] The formulation can also be extended to incorporate the selection of route times from among a finite set of alternative times. Allowing the model to shift route schedules is likely to reduce temporal conflicts and to increase profits.

The model is vitally dependent on the accurate estimation of profits by route and by aircraft. This involves the estimation of route operating costs and gross revenues by aircraft. The estimation of operating costs should be straightforward. Estimation of gross revenues is considerably more difficult as it necessitates predicting ridership levels which depend on fares, booking limits for different fare classes,[1, 14] schedules, and decisions made by competitors. The development of improved revenue estimation models is an important area of related research.

Finally, by exercising the model with alternative destinations, routes, fleet sizes and fleet compositions, the model may be used to assist planners in selecting routes or markets for expansion and in making fleet size and mix decisions. These too are important areas for additional research.

ACKNOWLEDGMENTS

THIS RESEARCH was supported by a National Science Foundation Presidential Young Investigator Award, grant 8352298 and a matching grant from the General Motors Corporation along with a grant from the United Parcel Service Corporation. This support is gratefully acknowledged. The authors also acknowledge the three anonymous referees and the editor-in-chief whose suggestions significantly improved the paper.

REFERENCES

1. P. BELOBABA, "Airline Yield Management: An Overview of Seat Inventory Control," *Trans. Sci.* **21,** 63–73 (1987).
2. H. CROWDER, "Computational Improvements for Subgradient Optimization," *Symp. Math.* **19,** 357–372 (1976).
3. M. S. DASKIN AND N. PANAYOTOPOULOS, "A Lagrangian Relaxation Approach to Assigning Aircraft to Routes in Hub and Spoke Networks," Transportation Center Working Paper TP-DAS-87.05, Northwestern University, Evanston, Ill., 1987.
4. M. M. ETSCHMAIER AND D. F. X. MATHAISEL, "Airline Scheduling: An Overview," *Trans. Sci.* **19,** 127–138 (1985).
5. M. L. FISHER, " The Lagrangian Relaxation Method for Solving Integer Programming Problems, *Mgmt. Sci.* **27,** 1–18 (1981).
6. M. L. FISHER, "An Applications Oriented Guide to Lagrangian Relaxation," *Interfaces* **15,** 10–21 (1985).
7. C. R. GLASSEY AND M. MIZRACH, "A Decision Support System for Assigning Classes to Rooms, *Interfaces* **16,** 92–100 (1986).
8. W. JUNGINGER, "Timetabling in Germany—A Survey," *Interfaces* **16,** 66–74 (1986).
9. A. LEVIN, "Scheduling and Fleet Routing Models for Transportation Systems," *Trans. Sci.* **5,** 232–255 (1971).
10. S. LIN AND B. W. KERNIGHAN, "An Effective Heuristic Algorithm for the Travelling Salesman Problem, *Op. Res.* **21,** 498–516 (1973).
11. N. PANAYOTOPOULOS, "An Application of Lagrangian Relaxation to an Aircraft Assignment Problem," M. S. thesis, Department of Civil Engineering, Northwestern University, Evanston, Ill. 1987.
12. L. T. PHILIPS, "Air Carrier Activity at Major Hub Airports and Changing Interline Practices in the United States' Airline Industry," *Trans. Res.* **21A,** 215–222 (1987).
13. G. SCHMIDT AND T. STROHLEIN, "Timetable Construction—An Annotated Bibliography," *Comput. J.* **23,** 307–316 (1980).
14. R. W. SIMPSON, "Setting Optimal Booking Levels For Flight Segments with Multi-Class, Multi-Market Traffic," in *Proceedings of the Twenty-fifth Annual AGIFORS Symposium,* pp. 263–279, Hikone, Japan, 1985.
15. F. SOUMIS, J. A. FERLAND AND J.-M. ROUSSEAU, "A Model for Large-Scale Aircraft Routing and Scheduling Problems," *Trans. Res.* **14B,** 191–201 (1980).
16. A. TRIPATHY, "School Timetabling—A Case in Large Binary Integer Linear Programming," *Mgmt. Sci.* **30,** 1473–1489 (1984).

(Received, July 1987; revisions received July 1988, October 1988; accepted October 1988)