Integrated Airline Planning

Rivi Sandhu (sandhu@uiuc.edu)
Diego Klabjan (klabjan@uiuc.edu)
Department of Mechanical and Industrial Engineering
University of Illinois at Urbana-Champaign
Urbana, IL

Abstract

The tactical planning process of an airline is typically decomposed into several stages among which fleeting, aircraft routing, and crew pairing form the core. In such a decomposed and sequential approach the output of fleeting forms the input to aircraft routing and crew pairing. In turn, the output to aircraft routing is part of the input to crew pairing. Due to the decomposition, the resulting solution is often suboptimal. We propose a model that integrates these three stages and thus yields a simultaneous solution to all three problems. We design two solution methodologies to solve the model. One is based on a combination of Lagrangian relaxation and column generation while the other one is the Benders decomposition approach. We give computational experiments for a variety of instances obtained by a major carrier.

1 Introduction

Airline business processes related to tactical planning consist of schedule planning, fleeting, aircraft rotation, and crew pairing (see e.g. Klabjan (2003a)). In *schedule planning*, a set of flights with specific departure and arrival times is constructed. Next is *fleeting*, which assigns an equipment type (such as Airbus 320, Boeing 737-500, etc.) to each individual flight. The objective of the *fleet assignment model* (FAM) is to maximize profit subject to the number of available aircraft and other operational constraints. The problems that follow decompose based on the fleeting solution, i.e. there is a separate problem for every equipment type. The *aircraft rotation* problem or *aircraft routing* is to find a set of generic aircraft routes that satisfy maintenance requirements. The *crew pairing* optimization follows. In crew pairing a set of crew itineraries or *pairings* is constructed. The goal is to minimize the crew cost and each flight must be covered by exactly one pairing. A few weeks before the day of operations, the actual tail numbers are assigned to each flight and monthly crew rosters or bidlines are assigned to every individual crewmember. Throughout tactical planning, revenue management related processes match demand with supply, which is the sit capacity of a given equipment type.

At present, each one of these problems is modeled independently of the remaining problems. (The only exception is a very recent research work focusing on combining selected two consecutive stages.) The models are solved sequentially and the output of one stage is the input to the next stage. Clearly there is interdependency between the various stages. For example, the fleeting solution decomposes the problems that follow by equipment type. The aircraft rotation and crew pairing problems are then solved over a subset of flights pertaining to a single equipment type. This is due to the fact that each cockpit crew is qualified to fly a particular set of equipment types. But this dependency of crews on equipment types is not captured in fleeting. In the aircraft routing stage decisions are made without considering the impact on the quality of the crew pairing solution. The interaction between these two stages is in the fact that a crew can have a connection shorter than a predefined number only if it stays on the same aircraft. Solutions obtained by using this sequential methodology can therefore be suboptimal. Ideally all the tactical planning problems should be solved as a single large-scale problem.

We present an integrated approach for these three most important stages. We propose a model that considers fleeting, aircraft rotation, and crew pairing simultaneously. Pairings are modeled explicitly and the rotation problem is captured by the plane count constraints from Klabjan et al. (2002). Thus when maintenance requirements are easily fulfilled due to the structure of the underlying flight network, our model completely integrates these three stages. We use two different solution methodologies to solve the integrated model. The first methodology entails solving the model by a combination of Lagrangian decomposition, Fisher (1985), and column generation. Column generation is used to price out favorable pairings based on the Lagrangian multipliers. By using Lagrangian relaxation some constraints are penalized and are moved to the objective function. These penalties are then used in pricing. The second solution methodology uses Benders decomposition, Benders (1962), on the model that relaxes the integrality of pairing variables. The traditional FAM model along with Benders cuts from the crew pairing LP relaxation forms the restricted master problem (RMP). Based on the incumbent fleeting solution in each iteration, for each fleet we solve the crew pairing LP relaxation and add a Benders cut to the RMP. The main contributions of our work are in the model itself and the novel solution methodologies for solving it. We believe this is the first work that models both fleeting and the crew pairing problems explicitly. The pricing problem is solved by constrained shortest path, where additional labels are required to capture different equipment types. Substantial profit improvements are obtained by our integrated approach.

In Section 2 we present a brief review of fleeting, aircraft routing, and crew pairing. The integrated model is presented in Section 3. Section 4 outlines both solution methodologies. Finally, in Section 5 we present the computational experiments.

2 Traditional Models

2.1 The Fleet Assignment Model

The fleeting problem is to find an assignment of equipment types to flights in a given flight schedule while maximizing profit, and subject to assignment constraints (each flight must be assigned to a fleet), flow balance (every aircraft that lands must take off), and plane count constraints (not to use more aircraft than there are available). The input is the flight schedule (obtained from the previous stage of schedule planning), the different equipment types and the available number of aircrafts for each equipment type. The objective function in FAM has two components – revenue and operating cost. The revenue is typically calculated based on the average fare for each leg. The operating cost consists of the costs associated by using a specific equipment type for a flight (fuel cost, landing fee, aircraft depreciation costs, etc.). We also incorporate the cost of spilled passengers for each fleet and flight in the operating cost. Since the model uses average fare for each leg to compute the revenue, such a model is also known as the *leg-based fleet assignment model* or traditional FAM.

The average fare for each leg is not representative of the actual revenue due to the multi leg passenger itineraries. Leg-based FAM solutions tend to be biased toward larger aircraft. This leads to the requirement of better revenue modeling practices. The *origin-destination fleet assignment model* (O-D FAM) models passenger revenues for each origin-destination itinerary rather than for each leg. The O-D FAM model is a combination of the leg-based FAM model and the passenger mix model. Given a fixed sit capacity, the passenger mix model decides the number of passengers that should be booked on a given O-D itinerary. The O-D FAM model is not as easy to solve as the traditional

FAM model. This is due to the fact that the possible number of passenger itineraries is very high and therefore the number of constraints and variables in the model is much larger. In our integrated model, we use the leg-based FAM model as opposed to O-D FAM. Incorporating O-D fleeting is part of the future research.

We next describe the FAM model, which is an important component of our integrated model. We first explain the underlying network required for model description. Consider a station o. An activity or event presents either a landing or a takeoff event. For the departure of flight l, let t_l be the departure time. For the arrival of flight l, let t_l be the arrival time plus the minimum aircraft turn time mt (typically around 30 minutes), called also the ready time. The *ready-time-space network* (RTN) has a node (o,s) for every station o and every activity s at this station. There is a *flight arc* between every departure and arrival event of the same flight. For every station o we order the activities based on t_l, i.e. $t_1 \leq t_2 \leq \cdots \leq t_n$, where n is the number of activities at the station (see the left figure in Figure 1). The network has a *ground arc* $g = ((o,s_i), (o,s_{i+1}))$ for i=1,2,..,n-1. In addition, there is a wraparound arc between the first and the last node of the time horizon. Let G denote the set of all ground arcs. The *ground arc time interval* is the time between the two activities that define the ground arc.

The FAM related variables are the binary fleet assignment variables x and the ground arc variables z. For each flight arc l and each equipment type f, x_{fl} is 1 if flight l is assigned to equipment type f. The nonnegative variable z_{gf} counts the number of planes on the ground of equipment type f during the ground arc g time interval. Let MD be a fixed time, which corresponds to a time with low activity at any station, e.g. 3 am. The FAM model reads

$$\max \sum_{f,l} r_{fl} x_{fl}$$

$$\sum_{f \in F} x_{fl} = 1 \qquad l \in L \tag{1}$$

$$\sum_{l \in O(v)} x_{fl} + z_{o(v)f} - \sum_{l \in I(v)} x_{fl} - z_{i(v)f} = 0 \qquad v \in V, f \in F \tag{2}$$

$$\sum_{l \in M} x_{fl} + \sum_{g \in W} z_{gf} \leq N_f \qquad f \in F \tag{3}$$

$$x \text{ binary}, z \geq 0,$$

where

$I(v)$	: set of flights to node v in RTN	L	: set of all flight arcs
$O(v)$	: set of flights from node v in RTN	F	: set of all fleets
M	: set of flights in the air at MD	V	: set of all nodes in RTN
N_f	: number of aircraft in fleet f	$i(v)$	: ground arc to node v in RTN
W	: set of ground arcs which contain MD	$o(v)$	: ground arc from node v in RTN
r_{fl}	: profit of assigning fleet f to leg l.		

This traditional FAM model is described in details in Hane et al. (1995). Details on the O-D FAM model and possible enhancements together with an extensive literature survey can be found in Klabjan (2003a).

2.2 Aircraft Routing

An *aircraft route* is a sequence of flights flown by the same aircraft or tail number. A *routing* or *rotation* is a set of aircraft routes, which partition all the flights in the schedule. The FAA imposes four types of maintenance checks. Three of them require a significant overhaul and are not captured in the aircraft routing problem. The so-called A checks must be performed every 3 to 4 days and typically have to be satisfied by a rotation. The aircraft routing problem also known as the *maintenance routing problem* consists of finding a maintenance feasible rotation and it must also ensures that the rotation does not use more than the available number of aircrafts. Since the fleeting solution decomposes the problem according to equipment types, the aircraft routing problem is solved for each fleet separately.

Since it is difficult to assign a cost to a given aircraft route, the problem is often treated only as a feasibility problem. There is a vast literature on the topic and we refer the reader to the survey work by Klabjan (2003a) for more detailed information. The problem is modeled either as a set partitioning problem, where columns correspond to maintenance feasible sequences of flights linking two maintenance stations, or as an Eulerian tour problem, which in turn leads to an asymmetric traveling salesman problem with side constraints. A third formulation is based on multicommodity flow approaches.

2.3 Crew Pairing

The crew pairing problem is to find a subset of pairings or crew itineraries that partition all the flights in the network while minimizing the crew cost. The input to this stage is the fleeting and the aircraft routes obtained in the previous stages. As with the aircraft routing problem, there is a separate crew pairing problem for each fleet. A pairing is a sequence of flights that satisfies several requirements. Clearly the destination station of a flight should be the same as the origin station of the next flight in the sequence. The origin station of the first flight should be the same as the destination station of the last flight in the sequence and it should be a *crew base*. A crew base is a designated station where crews are stationed. A pairing is composed of several *duties*, where a duty is a working day of a crew and is made up of a sequence of flights. A duty is subject to a number of FAA and union rules. Some of the duty legality rules are the maximum and minimum sit connection times between two consecutive flights, the maximum flying and elapsed time in a duty and many more. The cost of a duty is the maximum of three quantities: the flying time, a fraction of the elapsed time and the minimum guaranteed pay.

A pairing must also satisfy a large number of regulatory rules like the minimum and maximum rest time between two consecutive duties, the elapsed time, upper bound on the flying time, the FAR rules of the FAA, etc. The cost of a pairing is also the maximum of three terms: the sum of the duty costs, a fraction of the elapsed time and the minimum guaranteed pay. The problem is usually modeled as a set partitioning problem with side constraints. The decision variables y_p are equal to 1 if pairing p is part of the solution and 0 otherwise. The cost of a given pairing p is denoted by c_p. The resulting model reads

$$
\begin{aligned}
&\min \sum_{p} c_p y_p \\
&\sum_{l \in p} y_p = 1 \qquad \text{for each flight } l \\
&y_p \text{ binary.}
\end{aligned}
\tag{4}
$$

The crew pairing problem is difficult to solve due to the following two reasons. The number of pairings and thus variables is in the order of billions even for a medium size fleet of 200 flights. Also, the calculation of the cost of a pairing is very complex and a large number of complicated rules need to be taken into account while generating pairings. Due to the large number of variables, delayed column generation is employed. The pricing problem is traditionally solved as a constrained shortest path problem. Approaches based on finding the k^{th} shortest path or depth-first search enumeration of pairings on a network have also been proposed. The underlying network, which is used to generate pairings, is detailed in Section 4.1.1. More details on the crew pairing problem and a survey of literature on crew scheduling is provided by Barnhart et al. (2003).

2.4 Literature review

Some recent attempts have been made to integrate the various stages in airline planning. All of these attempts either only integrate two consecutive stages or they capture the crew pairing problem only at a very high level. Barnhart et al. (1998c) present a model that to some extent integrates FAM and crew pairing. They consider duties and not pairings in long-haul operations. The drawback of their model is that due to the large number of constraints, it is hard to solve. Barnhart et al. (1998a) propose a model for integrating FAM and the aircraft routing problem by using strings of flights. Desaulniers et al. (1997) present a model for the integration of FAM and time windows. Their model is a set partitioning model with side constraints. Rexing (1998) also presents a model that integrates FAM and time windows.

Cordeau et al. (2000) and Mercier et al. (2003) propose a model, which fully integrates the crew pairing and aircraft routing stages. They solve the model by a combination of branch-and-price and Benders decomposition. Integration of aircraft routing and crew pairing is also discussed in Cohn and Barnhart (2003), where each feasible routing is modeled as a column. Klabjan et al. (2002) present a model for partial integration of aircraft routing and crew scheduling. They solve the crew pairing problem before the aircraft routing problem by adding additional constraints to the crew pairing problem, which ensure feasibility of the aircraft routing problem in terms of the number of aircraft.

Similar integration efforts have been undertaken in mass transit scheduling. Haase et al. (2001) present a model, which minimizes both the crew costs and the number of vehicles. They solve the underlying set partitioning model with side constraints using a branch-and-cut-and-price algorithm. Freling et al. (2003) and Freling (1997) propose models that integrate vehicle and crew scheduling in a single depot environment.

3 *The Model*

The primary costs in an airline's operations are the operating costs from fleeting and the crew cost. The integration of FAM and crew pairing stages could thus potentially yield much lower crew costs and hence better profitability. But an integration of just FAM and crew pairing stages implies that the crew pairing problem is solved prior to the aircraft routing problem. The solution obtained from such an integrated model could potentially cause the maintenance routing problem to be infeasible in terms of maintenance and the plane count. In our integrated model, we first integrate the FAM and crew pairing problems. To prevent infeasible routings, we add additional constraints, which provide the necessary conditions for the aircraft routing problem to be plane count feasible. These necessary conditions ensure that the pairings do not violate the plane count, i.e. number of aircrafts on the ground at any point of time.

We assume that a daily flight schedule is given, i.e. every flight repeats every day of the week. The daily approximation is reasonable since for many carriers the majority of the flights have this property and the same approximation is used in the current sequential approach. We first describe a complete integration of the crew pairing and FAM problems. We integrate these models by enforcing that a pairing is assigned to a fleet if and only if all the flights, which constitute the pairing, are assigned to the same fleet. This implies assigning pairings to fleets, which requires expanding the crew pairing variables. We then link the new crew pairing variables with the fleet assignment variables. For fleeting, we use the aforementioned variables and constraints from the traditional FAM model described by (1)-(3). In order to capture pairings within fleeting, we modify the pairing variables to y_{fp}, where p is a pairing covering any subset of flights among all the flights in the schedule and f is a fleet index. y_{fp} is 1 if pairing p is assigned to fleet f. The fleet-pairing linking constraints must model that a pairing p assigned to fleet f covers flight l only if l is assigned to the same fleet f (determined by the corresponding x variable). We now detail the incorporation of the aircraft routing constraints in the integrated FAM and crew pairing model.

In traditional crew pairing, pairings are generated based on the minimum sit connection time *ms*, unless the crew follows the aircraft turn. We use the approach from Klabjan et al. (2002) to incorporate routing constraints. In the absence of maintenance constraints, they show that to integrate aircraft rotation and crew pairing it suffices to add the so-called plane count constraints to the set partitioning formulation of the crew pairing problem. Suppose we consider pairings with the minimum sit time equal to *mt*. Pairings from a solution with a connection shorter than *ms* (but by definition larger than *mt*) imply a plane turn. Such connections are called *forced turns*. A set of forced turns can be extended into a plane count feasible rotation if and only if the number of planes on the ground at any time imposed by the forced turns does not exceed the plane count obtained from the corresponding ground arc value in the FAM solution. The plane count constraints model this relationship. In order to embed this into our integrated model, we have to observe that the ground arc value now corresponds to a decision variable and is not a fixed value. A technical difficulty is the fact that the RTN uses ready times whereas pairings are based on the actual arrival times.

Let the *actual-time-space network* (ATN) be defined in the same way as RTN except that for each arrival of flight l, t_l is the actual arrival time (l_a) of leg l. The set of all ground arcs in ATN is denoted by G' and the ground arcs are denoted by g'. *Essential ground arcs* are those ground arcs, which are defined by a departure followed by an arrival. The plane count constraints associated with non-essential ground arcs are redundant (for proof see Klabjan et al. (2002)). Let $E(G') \subset G'$ be the set of essential ground arcs. Figure 1 details the difference between the RTN and ATN networks.

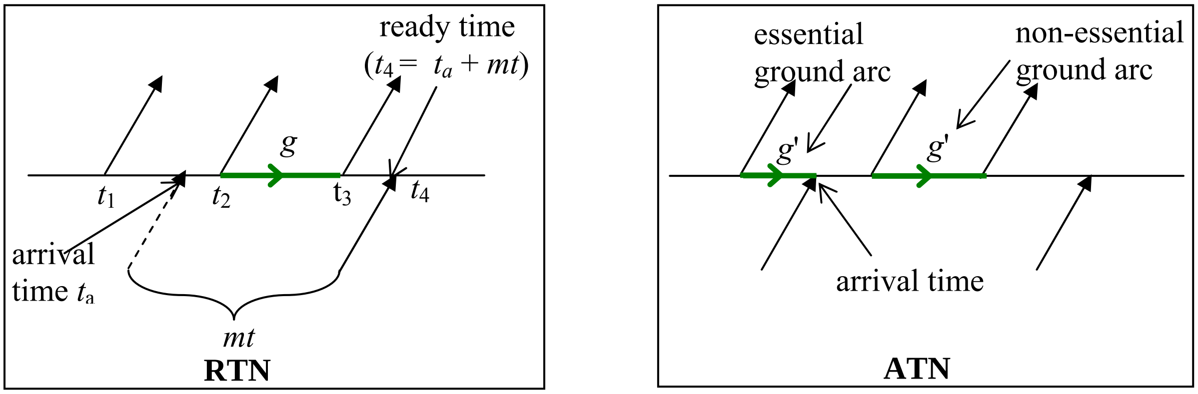

Figure 1: *Comparison between RTN and ATN*

Let P be the set of all pairings covering any subset of legs in the flight schedule with minimum sit connection time equal to mt. Let $P_{g'}$ be the set of pairings, which have a forced turn that includes g'. The plane count constraints from Klabjan et al. (2002) read

$$\sum_{p \in P_{g'}} y_p \leq b_{g'} \quad g' \in E(G'),$$

where, $b_{g'}$ is the number of aircraft on the ground in the time interval defined by ground arc g'.

It remains to be seen how to convert these constraints into a fleet based setting, i.e. we have to link $b_{g'}$ with the ground arc variables z. For each $g' \in E(G')$, there exists a corresponding ground arc $m_{g'}$ in RTN. Note that g' is defined by the arrival of a leg l. This ground arc $m_{g'}$ is defined by the ready time of leg l and an earlier activity. Let $dep(g')$ be the set of all flights that depart in the time interval $[l_a, l_a + mt]$. We replace the right hand side of the plane count constraints by

$$z_{m_{g'}f} + \sum_{l \in dep(g')} x_{fl}.$$

This expression states that the value of the ground arc corresponding to an arrival event of leg l in ATN is equal to the ground arc value of the same flight in RTN plus the number of departures in the set $dep(g')$. This correspondence between ground arc values in ATN and RTN is shown in Figure 2. The statement can easily be checked.

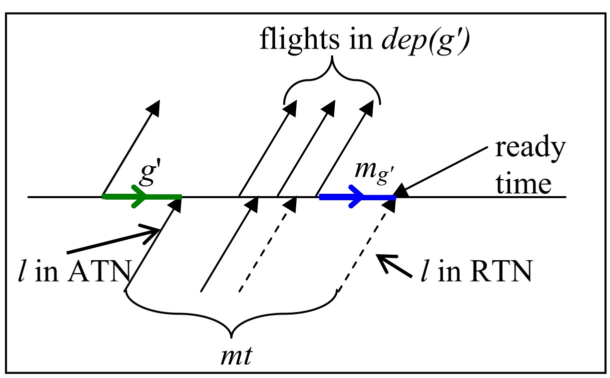

Figure 2: *Ground arc conversion*

The integrated model reads

$$\max \sum_{f,l} r_{fl} x_{fl} - \sum_{f,p} c_p y_{fp}$$

$$\sum_{f \in F} x_{fl} = 1 \qquad l \in L \tag{5}$$

$$\sum_{l \in O(v)} x_{fl} + z_{o(v)f} - \sum_{l \in I(v)} x_{fl} - z_{i(v)f} = 0 \qquad v \in V,\ f \in F \tag{6}$$

$$\sum_{l \in M} x_{fl} + \sum_{g \in W} z_{gf} \leq N_f \qquad f \in F \tag{7}$$

$$\sum_{l \in p} y_{fp} = x_{fl} \qquad l \in L,\ f \in F \tag{8}$$

$$\sum_{p \in P(g')} y_{fp} \leq z_{m_{g'}f} + \sum_{l \in dep(g')} x_{fl} \qquad g' \in E(G'),\ f \in F, \tag{9}$$

$$y \text{ binary}, x \text{ binary}, \mathrm{z} \geq 0.$$

Constraints (5)-(7) are the standard FAM constraints, (8) ensure that a pairing is assigned to a fleet if and only if all the legs in the pairing are assigned to the same fleet, and (9) are the plane count constraints.

Several extensions and enhancements can be easily incorporated. We mention a few of them. If plane turn times depend on the fleet (as is usually the case in practice), then all we have to do is to add dependency on f to g' and $dep(g')$ in (9). The corresponding RTN network needs to be adjusted accordingly. In practice, each fleet uses only a subset of crew bases. This is very easy to accommodate by changing the summation range in (8). Suppose that in each fleet f the number of crews must be in the range $[l_f, u_f]$, thus limiting the manpower. This is easily captured by adding

$$l_f \leq \sum_p y_{fp} \leq u_f \qquad f \in F$$

to the model. Other variations of the manpower constraints can be treated in a similar way.

4 Solution Methodologies

The integrated model is too large to be solved by standard optimization software packages even for small instances. Our approach is to either decompose the problem into smaller problems, which can be easily solved or to consider only a subset of pairings at a time. The reduced problem with only constraints (5)-(7) is the traditional FAM problem, which is relatively easy to solve. We propose two different methodologies.

In the first approach we use Lagrangian relaxation over a small subset of pairings to obtain Lagrangian multipliers. These multipliers are then used to price out new favorable pairings. In the Lagrangian relaxation step, we relax constraints (8) and (9) and then solve the resulting problem over a subset of pairings by using Lagrangian relaxation. The second methodology consists of relaxing the integrality requirements on the pairing variables. This problem can now be solved by Benders decomposition, wherein the crew pairing problem is solved as an LP. The information from the dual of the LP is used to form Benders cuts, which are then added to the FAM model. In Section 4.1 we describe the solution methodology using Lagrangian relaxation with column generation. A detailed discussion about column generation is postponed to Section 4.1.1. Section 4.2 details the solution methodology using Benders decomposition.

4.1 Lagrangian relaxation with column generation

Column generation is a methodology used to solve large-scale linear programs. In each iteration a smaller, reduced problem (the restricted master problem) with only a subset of columns/variables is solved. The dual information of the reduced problem is used to generate a new set of columns, which are added to the problem. The problem of generating new columns is referred to as the *pricing problem.*

Traditionally, column generation is used along with the LP relaxations in the branch-and-price framework (Barnhart et al. (1998b)). We propose a column generation type approach combined with Lagrangian relaxation. Lagrangian relaxation is used to generate multipliers that are equivalent to the dual vectors. First we approximate the model by changing the partitioning requirement (8) into covering constraints. We then relax constraints (8) and (9) and associate nonnegative Lagrangian multipliers λ_1, λ_2, λ_3,…with constraints (4) and μ_1, μ_2, μ_3,… with constraints (5). The restricted master problem reads

$$\phi(\lambda,\mu) = \max \sum_f \sum_l (r_{fl} + \lambda_{fl} - \sum_{g'} \sum_{l \in dep(g')} \mu_{fg'}) x_{fl} - \sum_f \sum_p (c_p - \sum_{l \in p} \lambda_{fl} + \sum_{g'} \sum_{p \in P(g')} \mu_{fg'}) y_{fp} - \sum_f \sum_{g'} \mu_{fg'} z_{fg'}$$

subject to FAM constraints (5),(6), and (7) and only over a subset of pairings.

The Lagrangian dual problem $\min_{\lambda \geq o, \mu \geq 0} \phi(\lambda,\mu)$ is solved by the subgradient algorithm. The obtained Lagrangian multipliers λ and μ are then used to find new pairings, which are added to the RMP. The RMP is the traditional FAM model with a different objective function which contains the lagrangian multipliers. The solution methodology is next detailed step-wise.

1. *Initialization*: Solve FAM and then next solve the crew pairing problem with the plane count constraints. We use the fleeting obtained from FAM to solve the crew pairing problem with the plane count constraints. The pairings obtained from crew pairing are used as the initial set of columns in the restricted master problem. The initial set of columns and this initial fleeting are feasible to integrated model since the pairings generated cover all the flights in the schedule and they do not violate the plane count constraints.
2. *Setting the lower bound*: We add the initial set of columns obtained in step 1 to the RMP. The objective value (profit in the integrated model) of the initial solution, denoted by *LB*, is equal to the objective value of the Lagrangian dual over this initial set of columns. This follows directly from the fact that the LP relaxation of the RMP in this case yields an integer solution. *LB* is clearly a lower bound on the optimal value of the integrated model as well as of the initial Lagrangian dual problem.
3. *Computing Lagrangian multipliers* (major iteration): We use subgradient optimization to solve the Lagrangian dual problem with the current set of columns. We use the step updating heuristic proposed by Caprara et al. (1999). We briefly describe the subgradients and the step size formula that we use.

   After changing the set partitioning constraint to a set covering constraint, constraints (8) read

$$\sum_{l \in p} y_{fp} \geq x_{fl} \qquad l \in L, f \in F.$$

Let x, y be an optimal solution to the RMP. The subgradient vector $s(\lambda)$ associated with a given λ is defined as

$$s_{l,f}(\lambda) = x_{fl} - \sum_{l \in p} y_{fp} \qquad l \in L,\ f \in F.$$

Similarly, the subgradient vector $s(\mu)$ associated with a given μ is

$$s_{g'f}(\mu) = \sum_{p \in P(g')} y_{fp} - z_{m_{g'}f} + \sum_{l \in dep(g')} x_{fl} \qquad g' \in E(G'),\ f \in F.$$

Let k be the iteration count within the subgradient algorithm. The step size along $s(\lambda^k)$, denoted by σ^k, is calculated as $\sigma^k = \theta\ (z(\lambda^k, \mu^k) - LB)/\|\ s(\lambda^k)\ \|^2$, where $\theta > 0$ is a parameter, which controls the step size along the subgradient direction. We change this parameter as required to increase or decrease the step size. For example, if the objective value does not change for k consecutive iterations, we reduce θ by half. We replace $s(\lambda^k)$ by $s(\mu^k)$ in the above formula and compute the step size τ^k along $s(\mu^k)$. Note that we deliberately do not use the same step size for λ and μ. The computational experiments have shown that such a choice yields better convergence.

Next we update the Lagrangian multipliers using the following formula

$$\lambda^{k+1} = \max\ \{\ \lambda^k + \sigma^k s(\lambda^k), 0\ \} \qquad l \in L,\ f \in F$$

$$\mu^{k+1} = \max\ \{\ \mu^k + \tau^k s(\mu^k), 0\ \} \qquad g' \in E(G'),\ f \in F\ ,$$

where the maximum is considered component-wise. We stop the subgradient algorithm either after a given number of iterations or if the norm of the subgradient becomes small.

4. *Pricing*: In this step we use the Lagrangian multipliers to price out favorable pairings using the constrained shortest path algorithm. Formally, we have to solve

$$\min_f \min_p \left(c_p - \sum_{l \in p} \lambda_{fl} + \sum_{g' : p \in P_{g'}} \mu_{fg'}\right) \tag{10}$$

This is detailed in Section 4.1.1. Let S be a small subset of pairings that either attain this minimum or are very close to it.

5. *Loop*: We first check the following termination criterion.
   - The above minimum (10) is nonnegative, i.e. no columns with negative reduced cost are obtained.
   - If $\phi(\lambda, \mu)$ does not change significantly for a given number of iterations.
   - If a predefined maximum number of major iterations have been completed.

   If any of these termination criteria is satisfied, we go to step 6; otherwise, we add the obtained pairings S from step 4 to the RMP and we go to step 3.
6. *Obtaining the final solution:* The final solution is obtained by using the fleeting produced by solving the Lagrangian dual in the last iteration. We use this fleeting to obtain aircraft routes and crew pairings using the traditional approach.

As described earlier, after fleeting is solved, the problem decomposes by fleet type. In the proposed integrated approach, the crew pairing problem for a given equipment type may be infeasible. This might happen either because we relax the partitioning constraints (8) to covering constraints or it might occur due to the nature of Lagrangian relaxation, i.e. a Lagrangian solution might not be feasible to the original problem. This was not observed in practice. Note that the traditional sequential approach suffers the same drawback. We believe that by using our integrated approach, which explicitly considers pairings and captures the dependency of pairings on fleets, the likelihood of producing a crew infeasible fleet is substantially lower.

4.1.1 Pricing

Pricing is the problem of generating pairings with the lowest reduced cost. There are two approaches to finding pairings with the least reduced cost. The first one is by enumerating all pairings and the second typically more efficient one is by using a variant of a shortest path algorithm. We use the latter approach to price out pairings based on (10). In this section, we first review the traditional constrained shortest path algorithm as it is applied to the crew pairing problem. We then show how to tailor this algorithm to solve (10).

There are two types of networks that can be built to solve the pairing generation problem: *flight network* or the *duty period network*. In the flight network each departure and arrival has an associated node in the network. A flight arc connects the departure node of a flight with the corresponding arrival node of the same flight. We add connection arcs between an arrival and a departure node at the same station subject to the constraints on the minimum connection time. We augment the network by adding a source s and a sink t. We connect source s to every departure node originating at a crew base. Similarly we connect all the arrival nodes from a crew base to sink t. The duty period network is similar except that we replace the flight arcs by duty periods and connection arcs correspond to legal rest connections. Although we can capture more pairing feasibility rules in the duty period network (all duty rules are embedded by definition), the duty period network is much larger than the flight network. Since we deal with the daily problem and to avoid cyclic networks, we replicate each flight several times until the maximum elapsed time of pairings is reached. For example, if pairings cannot exceed 5 days, then the network has 5 copies of each flight, each one offset in time by a day. The resulting network captures all pairings and it is acyclic.

It is clear that each pairing corresponds to an s-t path but an s-t path might violate pairing feasibility rules. In order to circumvent this, to find a favorable pairing the constrained shortest path algorithm must be employed. In such an algorithm, a label is maintained for each feasibility and cost resource. The latter are required if the cost of a pairing is non linear. Examples of labels are those corresponding to the maximum number of duties, the maximum elapsed time, the sum of the duty costs, etc. In addition, to capture the dual prices, an additional label is required. Each s-i path is represented as a vector consisting of the values of all resources. Thus every node i contains a set of label vectors.

In the constrained shortest path algorithm a node is selected and then all of its label vectors and adjacent nodes are scanned. For each label vector v and each neighbor node j a new label vector is formed by updating v using the underlying arc, and the new label vector is inserted at j. If a label vector is dominated by another one, it is discarded. For additional details on constrained shortest path see for example Desaulniers et al. (1998)**.**

For the integrated approach pricing (10), we can reduce the problem to the constrained shortest path problem. Assume that to solve the traditional crew pairing pricing problem we require k labels and let the corresponding label vector at node i be denoted by u. Furthermore, assume that it encodes a path p (partial pairing) from s to i. Since the Lagrangian multipliers contribute to the calculation of the reduced cost (10) we need to augment the underlying network and incorporate these values. We have $|F|$ Lagrangian multipliers for each flight and each ground arc. Therefore we need to add $|F|$ new labels to the existing labels. Thus the number of required labels becomes $k+|F|$. The new label vector becomes

$$v = (u, -\sum_{l \in p} \lambda_{1l} + \sum_{g' : p \in P_{g'}} \mu_{1g'}, -\sum_{l \in p} \lambda_{2l} + \sum_{g' : p \in P_{g'}} \mu_{2g'}, \cdots, -\sum_{l \in p} \lambda_{|F|l} + \sum_{g' : p \in P_{g'}} \mu_{|F|g'}).$$

It is easy to see that if $v \le \bar{v}$, then the path corresponding to $\bar{v}$ can be discarded and therefore this label vector can be removed.

If the duty period network is used, then for each duty d we precompute

$$(-\sum_{l \in d} \lambda_{1l} + \sum_{g' : d \in P_{g'}} \mu_{1g'}, -\sum_{l \in d} \lambda_{2l} + \sum_{g' : d \in P_{g'}} \mu_{2g'}, \cdots, -\sum_{l \in d} \lambda_{|F|l} + \sum_{g' : d \in P_{g'}} \mu_{|F|g'}).$$

These are then used as node weights in the network and the labels are updated upon scanning of duty d by adding this vector to the last $|F|$ coordinates of the label vector.

If the flight network is used, then the treatment is slightly different. For each flight arc l in the network we first form the vector $\alpha(l) = (-\lambda_{1l}, -\lambda_{2l}, \cdots, -\lambda_{|F|l})$. To handle ground arcs, consider a connection arc *ca* corresponding to a sit connection between an arrival of a flight and a departure of a different flight. Let the notation $g' \in ca$ mean that ground arc g' is included in connection *ca*. Next we form the vector

$$\beta(ca) = (\sum_{g' : g' \in ca} \mu_{1g'}, \sum_{g' : g' \in ca} \mu_{2g'}, \cdots, \sum_{g' : g'' \in ca} \mu_{|F|g'}).$$

An example is given In Figure 3, where the dotted lines show the actual flights, which constitute the essential ground arc. The solid lines are the flight arcs.

If we treat label vector v by scanning flight arc l, then the new label vector $\bar{v}$ is given by

$$\bar{v} = (\bar{u}, v_{k+1} + \alpha(l)_1, v_{k+2} + \alpha(l)_2, \cdots, v_{k+|F|} + \alpha(l)_{|F|}),$$

where $\bar{u}$ is the treatment of the standard crew pairing resources. If we are scanning a connection arc *ca* that corresponds to a sit connection, then the updated label vector is given by

$$\bar{v} = (\bar{u}, v_{k+1} + \beta(l)_1, v_{k+2} + \beta(l)_2, \cdots, v_{k+|F|} + \beta(l)_{|F|}).$$

If the connection arc corresponds to an overnight connection, then we treat only u.

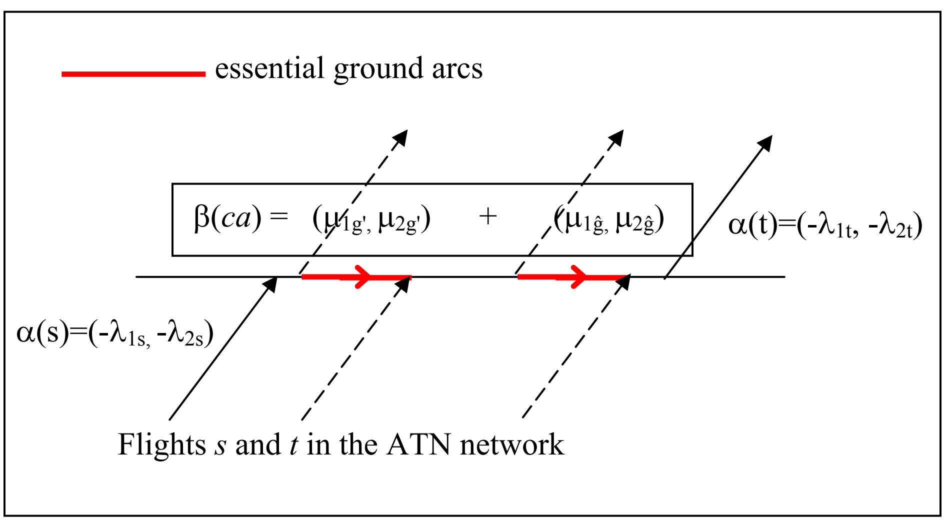

Figure 3: *Connection arc weights for* $|F| = 2$

In our computational experiments we solve (10) based on the duty period network.

4.2 Benders Decomposition

In this approach we relax the integrality of the pairing variables while maintaining the integrality on the fleet assignment variables. Clearly this yields a relaxation of the integrated model. By preserving the integrality of the fleet assignment variables, we do not relax the important revenue based decisions. In essence, we comply with the hierarchy of the current sequential decision making process.

The resulting model is now suited for the Benders decomposition approach. The restricted master problem consists of the FAM constraints (5)-(7) and the Benders cuts. Given a solution to the RMP, i.e. a feasible fleeting, the subproblem then decomposes into the LP relaxations of the crew pairing problems with plane count constraints. The subproblem clearly is decomposed by fleet and this substantially reduces the computational burden. More precisely, in each iteration we solve the RMP and then use the underlying fleeting to solve the LP relaxations of the crew pairing problems with embedded plane count constraints. By using the dual values of these LP relaxations, we add the Benders cut. If an LP relaxation is infeasible, then a Benders feasibility cut is added.

Next we elaborate on the most significant steps. For a given fleet f, we define L_f as the subset of flights that have been assigned to fleet f after solving the RMP, i.e. given a fleeting decision variables x. The LP subproblem for fleet f then reads

$$\min \sum_{p} c_p y_p$$

$$\sum_{l \in p} y_p = 1 \qquad l \in L_f \tag{11}$$

$$\sum_{p \in P(g')} y_p \leq b_{g'} \qquad g' \in E(G'_f) \tag{12}$$

$$y_p \geq 0.$$

Here $E(G'_f)$ denotes the set of the essential ground arcs with respect to the flights in L_f and $b_{g'} = z_{m_{g'}} + \sum_{l \in dep(g')} x_{fl}$ is the ground arc value.

Let the duals for (11) be represented by θ and the duals for (12) by Π, whenever the LP is feasible. If the LP is infeasible, let (α, β) be an extreme ray with

$$\sum_{i \in L_f} \alpha_{fl} + \sum_{g' \in E(G'_f)} b_{g'} \beta_{g'f} > 0, \tag{13}$$

where α corresponds to (11) and β to (12). Let η represent the upper bound on the crew cost in the relaxed integrated model. We index the Benders cut by k, where k is the iteration count. The RMP at the beginning of a new iteration reads

$$\max \sum_{f,l} r_{fl} x_{fl} - \eta$$

subject to FAM constraints (5), (6) and (7),

$$\eta \geq \sum_f \sum_l \left[\theta_{fl}^k + \sum_{g': l \in dep(g')} \Pi_{g'f}^k \right] x_{fl} + \sum_f \sum_{g'} \Pi_{g'f}^k z_{m_{g'}f} \qquad k \in K \tag{14}$$

$$\sum_l \left[\alpha_{f_k l}^k + \sum_{g': l \in dep(g')} \beta_{g' f_k}^k \right] x_{f_k l} + \sum_{g'} \beta_{g' f_k}^k z_{m_{g'} f_k} \leq 0 \qquad k \in J \tag{15}$$

η unrestricted, x binary, $z \geq 0$.

Constraints (14) are the Benders cuts and (15) are the Benders feasibility cuts. Here f_k denotes an infeasible subproblem LP at iteration k.

The entire solution methodology is described stepwise as follows. Unlike the first solution methodology, we do not need to generate an initial feasible solution in this case.

1. *Solving the Restricted Master Problem:* Solve the restricted master problem and obtain a fleeting.
2. *Decompose based on fleeting:* Based on the solution of the RMP, for each fleet we generate legs and plane count information.
3. *Solve the crew pairing LP relaxations:* For a given fleet, we use the corresponding leg and plane count information obtained in step 2 to solve the LPs (11)-(12). Once the LPs have been solved, one of the following two cases may arise: the solution is optimal for each fleet or there is a fleet that yields an infeasible subproblem.
4. *Generate Benders cut:* If all the subproblems are feasible, we obtain the duals θ^k and Π^k for constraints (11) and (12) respectively for the current iteration k. Use these duals to add a new benders cut (14) and go to step 6.
5. *Generate a feasibility cut:* For each infeasible subproblem, we obtain an extreme ray satisfying (13). For each such infeasible fleet, we add the corresponding feasibility cut (15).
6. *Iterate:* If we reach the maximum number of iterations, we go to step 7. Otherwise we go to step 1.
7. *Obtaining the final solution*: We use the fleeting in the last iteration to generate aircraft routes and crew pairings using the traditional approach.

5 Computational Experiments

The computing environment consists of a cluster of 27 dual 900 MHz Itanium 2 processors running Red Hat 7.3 operating system and the gcc 3.2 development environment. For solving small LPs and the fleeting integer programming models we used CPLEX from ILOG Inc., version 8.1. We tested the integrated model on three data sets. Real world data from a major US carrier were used. The carrier has a heavy hub-and-spoke network structure with five crew bases and 8 hubs. Crew feasibility rules and cost function comply with airline rules. For discretionary purposes, the real profit numbers are fudged but the presented numbers show correct proportions and magnitudes. Due to the lack of data, we do not use manpower constraints. However, the minimum turn time is fleet dependent and each fleet uses only a specific subset of crew bases.

In order to make a more fair comparison between the proposed solution methodologies and the traditional sequential approach, in our implementation we have also modeled constraints to prevent crew double overnights (Clarke et al. (1994)). We add constraints to FAM using legal rest arcs and mid-day breakouts such that the obtained fleeting is more “crew friendly”.

Each crew pairing problem is solved by branch-and-price. The pricing problem (10) and the pricing involved in solving the crew pairing LP relaxations in the second approach is carried out by using the parallel constrained shortest path algorithm, Klabjan (2003b).

We control tractability in the following way. First FAM is solved over all fleets and flights. Next we pick a subset of fleets and the corresponding flights. The integrated model is then solved by considering only this subset of fleets and flights. In other words, instead of considering all fleets at once, we can consider subsets of fleets. The number of flights in each instance and the number of considered fleets are given in Table 1. Cases 2 and 3 are not full fleets, i.e. after obtaining a fleeting, we choose a subset of fleets and then from within each set of flights for a given fleet we further pick a subset of flights. They were generated as a proof of concept instances. Case 1 on the other hand consist of the entire set of flights for a subset of fleets. In this section we present the computational results for the three test cases. We first discuss all the benefits and then give a more detailed analysis of some of the test cases.

The increase in profits obtained by using the integrated approach as opposed to the sequential one is shown in Table 1. The unit is US dollar. Both approaches were stopped after approximately the same elapsed time, which is discussed later. Profit here is defined as the combined profit resulting from FAM and the crew pairing costs, i.e. for the integrated approach it corresponds to the objective value. In the last two columns, we show the increase in profit obtained by solving the integrated model by Lagrangian and Benders respectively. Except for case 3, Lagrangian relaxation approach outperforms the Benders decomposition. We stress that these are daily profit increases that produce on an average 50 million dollars of additional profit per year. In addition, we selected a subset of fleets just once. By repeating the process several times additional profit can be obtained.

	Test Case		**Increase in Profit ($)**	
	Fleets	***Flights***	***Lagrangian***	***Benders***
Case 1	2	524	160,000	150,000
Case 2	3	251	30,000	8,000
Case 3	2	205	22,000	33,000

Table 1: *Profits*

Table 2 details the breakup of the profit by revenue, operating cost for operating the flights and the crew cost. The breakup has been presented for both solution methodologies. (The abbreviation for Lagrangian is Lgr and for Benders is Bns.) All values are shown as percentage increase or decrease values with respect to the corresponding value obtained by using the traditional methodology. To ensure data confidentiality, we have presented the range of percentile instead of the actual percentile increase or decrease. For example, for case 1 solved by the Lagrangian approach the revenue decrease is lower than 2% with respect to the traditional approach. Most of the increased profit does not come from diminishing revenue in FAM but it comes from the reduced crew cost. The change in the operating cost is negligible. This is a desired property since for historical and cultural reasons the carriers are not willing to sacrifice too much on the revenue side (even though it would increase the total profit).

	Revenue %		Operating Cost %		Crew Costs %		Savings %	
	Lgr	***Bns***	***Lgr***	***Bns***	***Lgr***	***Bns***	***Lgr***	***Bns***
Case 1	[-1,0]	[-1,0]	[0,1]	[0,1]	[-11, -8]	[-10, -8]	[5,8]	[4,7]
Case 2	[-1,0]	[-1,0]	[0,1]	[0,1]	[-6, -3]	[-4, -1]	[1,4]	[0.4,1]
Case 3	[-2,0]	[-1,0]	[0,0.5]	[-1,0]	[-4, -2]	[-7, -5]	[1,4]	[2,5]

Table 2: *Breakup of profit*

We next study in more details cases 1 and 3 when solved using the Lagrangian approach. Figure 4 shows the improvement in the objective value after every major iteration (steps 3-5 in the algorithm) for the two test cases. Figure 4a shows the improvement for case 1 while Figure 4b shows the improvement for case 3. Obviously in initial iterations the improvements are minimal and then the objective value increases substantially. It is obvious that based on this trend it would be beneficial to perform additional iterations.

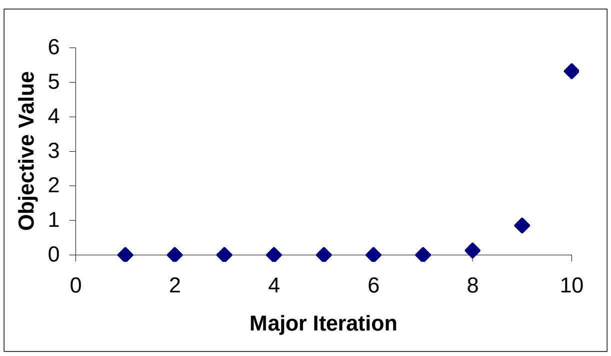

Figure 4a)**:** Case 1

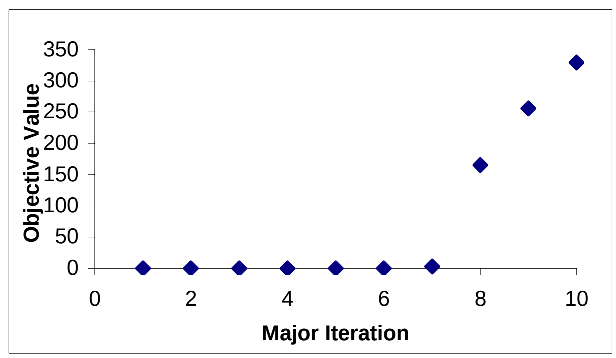

Figure 4b): Case 3

Figure 4: *Objective value improvements*

Within a given major iteration the objective values $\phi(\lambda,\mu)$ of the Lagrangian relaxation tend to decrease and eventually they converge towards the value of the Lagrangian dual. Figure 5 and Figure 6 show the decreasing trend of the Lagrangian relaxation objective values within a given major iteration for cases 3 and 1 respectively. Figure 5a and Figure 6a show the trend in a relatively early major iteration while Figure 5b and Figure 6b show the trend in a later major iteration. In few iterations this objective value significantly increases. We do not have an explanation for such a behavior. Fortunately it decreases back to the previous value relatively soon.

It also seems that the objective value is more stabilized in early major iterations than in the late iterations. Also with an increased data set size, the oscillation tends to increase. It is intuitive that convergence is slower when the data set size increases, which is apparent by comparing Figure 6 with Figure 5. Case 1 is larger then case 3.

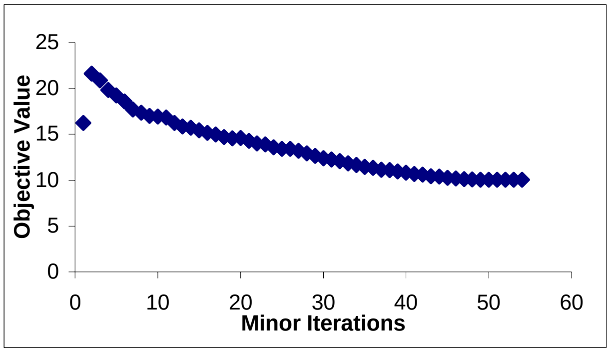

Figure 5a)**:** Early major iteration

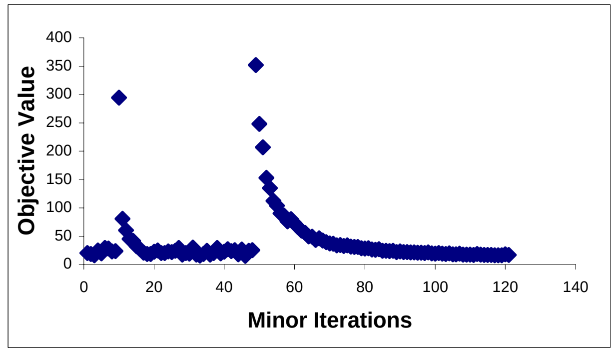

Figure 5b): Late major iteration

Figure 5: *Trend of the Lagrangian relaxation objective values for case 3*

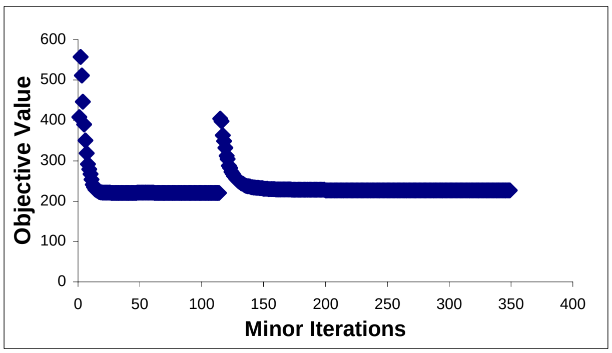

Figure 6a)**:** Early major iteration

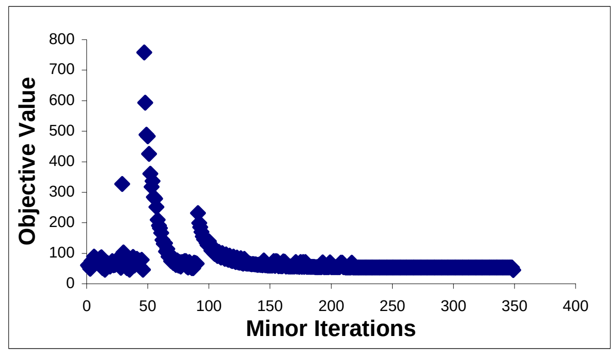

Figure 6b): Late major iteration

Figure 6: *Trend of the Lagrangian relaxation objective values for case 1*

Now we focus on the Benders decomposition. In Benders decomposition the objective value is non-increasing across iterations. The trends for cases 1 and case 3are shown in Figure 7. As we can see, due to degeneracy in subproblems, the objective value can be constant for some iterations. We did not experiment with the core approach from Magnanti and Wong [22] since it is non trivial to find a core point.

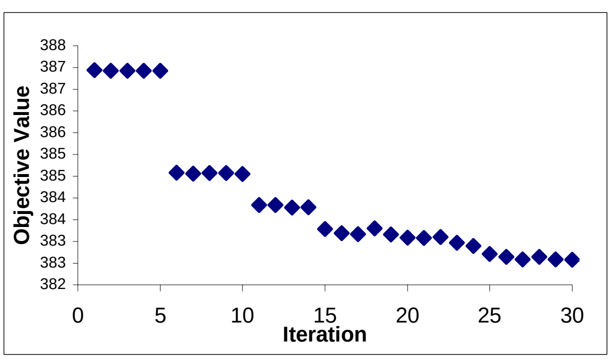

Figure 7a)**:** Case 2

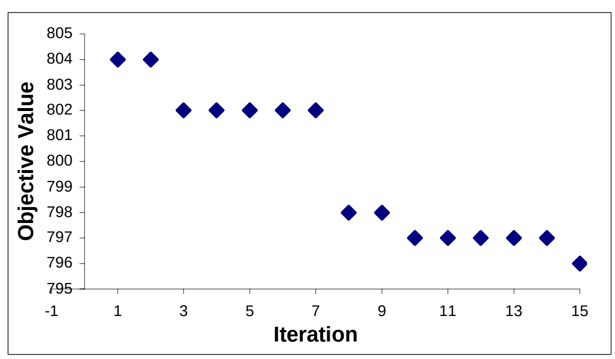

Figure 7b): Case 4

Figure 7: *Objective value in Benders decomposition*

The breakup of computation times obtained by solving the integrated model using the Lagrangian approach is shown in Table 3. The units of measurement are listed in each column. The Lgr column shows the time taken to solve one subgradient iteration. The "average" column for each of the FAM, Lgr and Pricing steps shows the average cumulative running time over all iterations. The percentile column for each of these steps shows the percentage of the total time that is spent in solving the step over all iterations. (We sum up the time taken for a given step in all minor and major iterations and calculate the % of the total time, which can be attributed to the given step.) The last column shows the total time taken to solve a given data set using the integrated model.

	Per Major Iteration	Per Minor Iteration	FAM		Lgr		Pricing		Total
			Average	%	Average	%	Average	%	
	sec	sec	sec		sec		sec		Hrs
Case 1	8,700	31	21	67	27	80	1,500	17	25
Case 2	6,600	50	37	70	49	89	600	9	19
Case 3	5,400	16	12	78	15	93	180	3	15

Table 3: *Computation times for the Lagrangian approach*

The running time for each major iteration increases significantly with an increasing number of legs. There is not a considerable increase in computational time with an increase in the number of fleets as shown by cases 2 and 3. The subgradient optimization accounts for more than three-fourths of the time. The most intensive step is solving the RMP in each iteration and all other operations per iteration are negligible. The reason is in the slow convergence of the subgradient algorithm. It is surprising that the pricing step is not computationally intensive. This is primarily due to the employment of the parallel algorithm.

The computation times for Benders are shown in Table 4. The "average" column for the RMP and LP steps shows the absolute value of the average over all iterations. The "%" column shows the percentile of the total time that is spent in solving a step over all iterations. The last column gives the overall running time. The increase in the running time to solve the RMP is striking. The traditional FAM used in the Lagrangian approach takes less than 1 minute. On the other hand, after adding Benders cuts (we never add more than 30 of them), this running times increases to almost an hour.

	Time per Iteration	RMP		LP subproblem		Total
		Average	**%**	**Average**	**%**	
	min	min		min		Hrs
Case 1	82	58	71	21	26	19
Case 2	66	57	88	7	11	15
Case 3	22	17	79	4	18	5

Table 4: *Computation times of Benders decomposition*

With an increase in the number of legs and fleets, the computational time using the Benders approach increases. The distribution of time across RMP and LP shows a significant trend. For a smaller test case, the RMP takes about three-fourth of the total time to solve the problem. But with an increase in the size of the data set, the time taken to solve the LP subproblem increases significantly. For a larger data set, half of the total time is spent in solving the crew pairing LP relaxation.

6 Future Directions

The convergence rate of the Lagrangian algorithm using the subgradient optimization is very slow. The bundle method is known to have a better convergence rate as compared to subgradient variants to solve Lagrangian relaxation. We see a potential improvement in the running time by using the bundle algorithm instead of the vanilla subgradient optimization.

On the algorithmic front, we plan to test an alternative algorithm, which is aligned with the existing one. We propose to directly solve the Lagrangian dual. In the current solution methodology, we consider only a subset of pairings in each iteration. We then generate columns and add them to the current subset. As opposed to this, in every subgradient or bundle iteration, we can price out all pairings to adjust the current iterate.

The revenue obtained from the current integrated model takes only the average fare per leg into account, i.e. fleeting is leg based. To overcome this problem and to be consistent with yield management practices, we can integrate O-D FAM, crew pairing, and aircraft routing. O-D FAM can be solved using Benders decomposition. The newly formed integrated model can be solved along the same lines as the current integrated model. Unfortunately we believe this still might be out of reach since the solution times for O-D fleeting are too high. An alternative might be to use clever high performance algorithms.

In our integrated model, although we present a complete integration of FAM and crew pairing problems, the integration with the aircraft routing model is only partial. Recall that out solution yields plane count feasible aircraft routes but not necessarily maintenance feasible routes. In addition, we do not capture the potential aircraft routing objective function. An integrated model, which fully integrates the aircraft routing model with the FAM and crew pairing models instead of just ensuring feasibility, is an interesting topic for future research.

7 Acknowledgments

We would like to thank Michael Clarke from Sabre Inc. for providing the data and valuable discussions. We are indebted to Professor Karsten Shwan from Georgia Institute of Technology for allowing us to use their computing resources. Special thanks to Professor Eva Lee from the Georgia Institute of Technology, Atlanta, GA for assistance with CPLEX licenses. In addition, we are thankful to ILOG Inc. for providing the licenses.

8 References

1. Barnhart, C., Cohn, A., Johnson, E., Klabjan, D., Nemhauser, G. and Vance, P. (2003), Crew scheduling. In Hall, R, editor, *Handbook of Transportation Science*, 517-560. Kluwer Scientific Publishers.
2. Barnhart, C., Boland, N., Clarke, L., Johnson, E., Nemhauser, G., and Shenoi, R. (1998a). Flight string models for aircraft, fleeting and routing. *Transportation Science*, **32**, 208-220.

3. Barnhart, C., Johnson, E., Nemhauser, G., Savelsbergh, M. and Vance, P. (1998b). Branch-and-Price: Column generation for solving huge integer programs. *Operations Research*, **46**, 316-329.
4. Barnhart, C., Lu, F., and Shenoi, R. (1998c). Integrated airline schedule planning. In Yu, G., editor, *Operations Research in the Airline Industry*, **9**, 384-403. Kluwer Academic Publishers.
5. Bélanger, N., Desaulniers, G., Desrosiers, J. and Soumis, F. (2003). Periodic airline fleet assignment with time windows, spacing constraints and time dependant revenues. To appear in *European Journal of Operational Research.*
6. Benders, J. (1962). Partitioning procedures for solving mixed-integer programming models. *Numerische Mathematik*, **4**, 238-252.
7. Caprara, A., Fischetti, M. and Toth, P. (1999). A heuristic algorithm for the set covering problem. *Operations Research,* **47**, 730-743.
8. Clarke, L., Hane, C., Johnson, E., and Nemhauser, G. (1994). Maintenance and crew considerations in fleet assignment. *Transportation Science*, **30**, 249-260.
9. Cohn, A., and Barnhart, C. (2001). Improving crew scheduling by incorporating key maintenance routing decisions. *Operations Research*, **51**, 387-396.
10. Cordeau, J., Desrosiers, J., Soumis, F, and Stojković, G. (2000). Benders decomposition for simultaneous aircraft routing and crew scheduling. *Transportation Science*, **35**, 375-388.
11. Desaulniers, G., Desrosiers, J., Dumas Y., Solomon, M., and Soumi, F. (1997). Daily aircraft routing and scheduling. *Management Science*, **43**, 841-855.
12. Desaulniers, G., Desrosiers, J., Ioachim, I., Solomon, M., and Soumis, F. (1998). A unified framework for deterministic constrained vehicle routing and crew scheduling problems. In Crainic, T. and Laporte, G., editors, *Fleet Management and Logistics*, 57-93. Kluwer Publishing Company.
13. Fisher, M. (1985). An applications oriented guide to Lagrangian relaxation. *Interfaces*, **15**, 10-21.
14. Freling, R. (1997). *Models and techniques for integrating vehicle and crew scheduling.* PhD thesis, Tinbergen Institute, Erasmus University Rotterdam.
15. Freling, R., Huisman, D., and Wagelmans, A. (2003). Models and algorithms for integration of vehicle and crew scheduling. *Journal of Scheduling*, **6**, 63-85.
16. Haase, K., Desaulniers, G., and Desrosiers, J. (2001). Simultaneous vehicle and crew scheduling in urban mass transit systems. *Transportation Science,* **35**, 286-303.
17. Hane, C., Barnhart, C., Johnson, E., Marsten, R., Nemhauser, G. and Sigismondi, G. (1995). The fleet assignment problem: Solving a large-scale integer program. *Mathematical Programming,* **70**, 211-232.
18. Klabjan, D. (2003a). Large-scale models in the airline industry. *Technical Report*. To appear in GERAD 25th Anniversary Volumes. Available at https://netfiles.uiuc.edu/klabjan/www**.**
19. Klabjan, D. (2003b). *Parallel constrained shortest path*. Presentation at ISMP 2003, Copenhagen, Denmark.
20. Klabjan, D., Johnson, E., Nemhauser, G., Gelman, E., and Ramaswamy S. (2001). Solving large airline crew scheduling problems: Random pairing generation and strong branching. *Computational Optimization and Applications*, **20**, 73-91.
21. Klabjan, D., Johnson, E., Nemhauser, G., Gelman, E., and Ramaswamy S. (2002). Airline crew scheduling with time windows and plane count constraints. *Transportation Science,* **36**, 337-348.

22. Magnanti, T., and Wong R. (1981). Accelerating Benders decomposition: Algorithmic enhancement and model selection criteria. *Operations Research*, **29**, 464-484.

23. Mercier, A., Cordeau, J. –F., and Soumis, F. (2003). A computational study of Benders decomposition for the integrated aircraft routing and crew scheduling problem. To appear in *Computers & Operations Research*.

24. Rexing, B. (1998). *Fleet assignment with time windows*. Master's thesis, Massachusetts Institute of Technology.