

Chapter 2

Airline Planning and Schedule Development

Timothy L. Jacobs, Laurie A. Garrow, Manoj Lohatepanont, Frank S. Koppelman, Gregory M. Coldren, and Hadi Purnomo

2.1 Introduction and Scope

Airlines have evolved over the past 70 years from simple contract mail carriers into sophisticated businesses. The current airline environment is very competitive and dynamic. Maintaining consistent profitability requires that appropriate trade-offs be made between the often competing objectives within planning, marketing and operations. Airlines have led other industries in the application of operations research and information technology to deal with these issues. The real-time solution of large-scale optimization models has played a significant role in shaping today's airline industry. This role will increase as the industry becomes more competitive and flight characteristics change due to the implementation of new technologies. Airline planning and scheduling represents an excellent example of the application of operations research and mathematical modeling to solve complex and real industry problems.

2.2 Overview of Airline Schedule Planning and Marketing

Planning and Marketing define an airline's products and determine how they will be sold. This is a continuous process which begins 5 or more years before a flight's departure and operates until the last passenger is boarded and the aircraft door is

T. L. Jacobs (✉) · L. A. Garrow · M. Lohatepanont · F. S. Koppelman ·
G. M. Coldren · H. Purnomo
e-mail: Tim.Jacobs@usairways.com

L. A. Garrow
e-mail: laurie.garrow@ce.gatech.edu

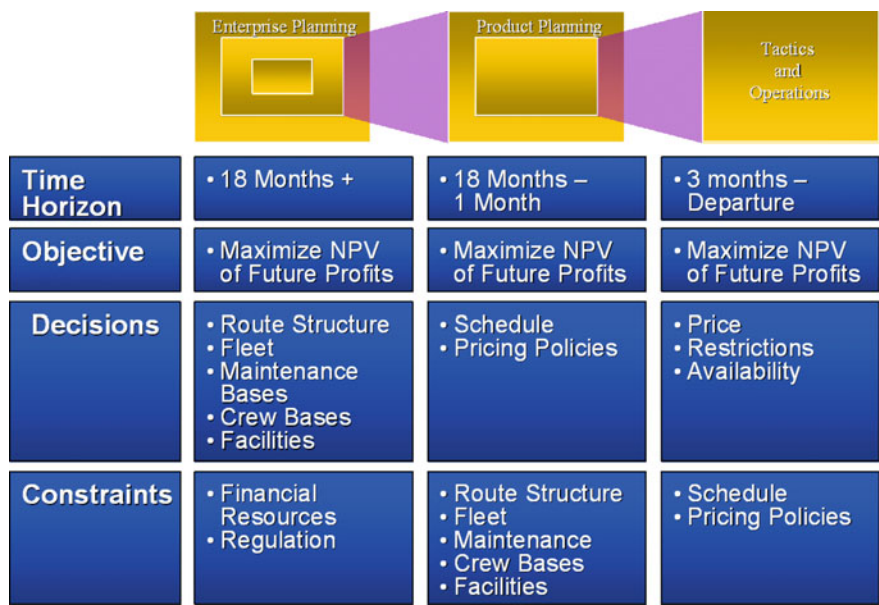


Fig. 2.1 Schedule planning and operation timeframe, objectives and constraints (Smith and Jacobs 1997)

closed. This process can be viewed as a series of overlapping sequential steps that include scheduling, marketing and distribution. This process requires an exchange of data and feedback between scheduling, pricing and revenue management and distribution. In addition, other considerations such as crew resources, maintenance and engineering and ground services help define the boundaries by which the airline schedule must operate and be managed (Fig. 2.1).

Scheduling determines where and when the airline will fly. Schedules are built to maximize long-term profitability. The revenue and cost associated with each schedule are based on very different views of the same information. Although the schedule is composed of individual flight legs between two cities, the airline’s product and revenues are based on passenger origin and destination (O&D) markets. An O&D market is defined by a passenger’s point of entry and exit from the airline system. The schedule is built to maximize its attractiveness to customers in a wide variety of O&D markets. The development of hub and spoke networks was based on providing maximum O&D coverage with a limited number of flight legs. The costs of operating the schedule depend on the flight legs, which drive the number and type of aircraft used. The schedule must consider the cost and availability of cabin and flight deck crews, as well as the requirement that aircraft cycle through maintenance bases at regular intervals. As a result, the schedule also determines the location and size of ground facilities, and the number and location of crew and maintenance bases.

Efficient schedules which match supply and demand are key to airline profitability. Profitable solutions require anticipation of general market conditions: the costs of capital, fuel and labor, as well as the level and nature of competition. Airlines address many scheduling issues (assigning aircraft and crews to flights, routing aircraft to maintenance bases) with large-scale combinatorial optimization techniques. The scale of today's airlines makes this increasingly difficult. For example, large U.S. domestic carriers operate more than 3,000 flights per day with 600 or more aircraft and can include more than 300 cities, serving over 10,000 unique O&D markets.

Marketing determines what specific products will be offered for sale and how many of each will be sold. The two primary components of airline marketing are pricing and yield management. Since deregulation of the U.S. domestic airline industry, both have evolved into very complex processes. Prior to deregulation, individual airlines served specific market segments. Scheduled carriers served the business traveler while charter carriers served the leisure market. Scheduled carriers flew with relatively low load factors but remained reasonably profitable due to the limited competition created by government regulation.

Just prior to deregulation, the scheduled carriers started to offer additional products to the leisure market segment to help fill some of the empty seats. There were two problems with this approach. First, airline revenue would have been severely diluted if the existing customer base switched from full fare to discounts. Restrictions were introduced to make discounts unappealing to existing business travelers (advance purchase, saturday night stay). Second, some flights were already full and discount sales would displace late booking, high value traffic. The yield management process was introduced shortly after deregulation to anticipate where discount sales would and would not be profitable. From this simple beginning, the pricing and yield management process has become very complex and dynamic. Their combined role is to help airlines fine tune demand and sales to meet the capacity provided by the schedule. Today, a U.S. domestic carrier's schedule can consist of up to 4 million fares. These fares and restrictions are adjusted frequently to match demand to supply; 100,000 fare changes per day is typical. As a result, revenue management departments must keep up with these changes. Every day the number of reservations available for sale is reviewed and adjusted in order to maximize total revenues on all future flights. This is another very large-scale optimization process that involves solving a highly stochastic and nonlinear optimization model. Controlling reservation availability for all future flights at a large U.S. carrier can represent a problem with approximately 500 million decision variables.

Distribution is the process of taking the airline products and putting them on the shelf for sale. The store front for the airline industry is primarily central reservation systems (CRS) and global distribution systems (GDS). A CRS allows an airline's reservation agents to book their own flights and fares. CRSs are relatively expensive to develop and maintain. Large airlines typically have their own CRS, while second and third tier carriers tend to rent space in another carrier's (often their competitor's) system. In the 1970s U.S. domestic carriers began to give travel agents access to their CRSs. This provided each airline with a method for selling products outside of the

reservations office. Through the 1980s the number of travel agents with direct access to the major CRSs (Sabre, Apollo-United Airlines, Galileo-British Airways, World Span) grew substantially. Each CRS became a significant distribution outlet not only for the originating airline but for other participating airlines. Today the successors to these CRSs have become global in that an agent hooked into Sabre or Amadeus has access to the schedules, fares and reservation availability for most of the world's airlines. This gives any airline immediate access to a very significant distribution process. For example, Sabre contains schedules for over 700 airlines and agents can book tickets directly on 350 airlines. Sabre is installed in over 29,000 travel agency locations and processes more than 4,900 messages per second.

2.3 Chapter Outline

This chapter will focus on the application of forecasting and operations research techniques to airline scheduling problems and provide a brief overview of how airlines use these techniques to develop and evaluate schedules and business decisions. [Section 2.2](#) of this chapter provides a brief overview of the forecasting process used to estimate passenger demand and determine the expected revenue, cost and profitability associated with a given schedule. This section will also provide insight as to how airlines use these techniques to evaluate incremental changes to market services such as frequency and/or aircraft assignment changes. [Section 2.3](#) provides an overview of the fleet process used by most network carriers. This section presents an introduction to the fleet assignment model (FAM) and some of the enhancements to the model better integrate the schedule development and fleet processes with both operational and revenue management aspects of the airline business process. These enhancements will include the incorporation of O&D passenger effects (O&D FAM) and the inclusion of high-level maintenance and engineering (M&E) opportunities into the schedule development and fleet process. [Section 2.4](#) will provide a high-level overview of the aircraft routing process and its impact on other business units within the company such as M&E and flight and ground crew resources. [Section 2.5](#) presents an overview of some new developments and directions in the operations research and forecasting and their application to the airline scheduling area. [Section 2.6](#) provides a full list of references noted within the chapter for further study.

2.4 Forecasting Aspects and Methodologies for Schedule Planning

This section includes three major components. First, an overview of the data and major components of network-planning models is presented. Next, the two major types of market share models based on the Quality of Service Index (QSI)

methodology or logit-based methodologies are reviewed. This is followed by a summary of the key experiences of a major U.S. airline that transitioned from using an itinerary choice model based on a QSI methodology to one based on a logit methodology. The section concludes with a discussion of one relatively new modeling technique: the use of continuous time-of-day functions (versus discrete time-of-day dummy variables or schedule delay functions). Part of the material in this section is reprinted from Garrow (2010, pp. 203–208, 228–229, 250) with permission from Ashgate Publishing. The material draws heavily on prior work from Coldren and Koppelman as well as information obtained via interviews with industry experts.

2.4.1 Introduction

Network-planning models (also called network-simulation or schedule profitability forecasting models) are used to forecast the profitability of airline schedules. These models support many important long- and intermediate-term decisions. For example, they aid airlines in performing merger and acquisition scenarios, route schedule analysis, code-share scenarios, minimum connection time studies, price-elasticity studies, hub location and hub buildup studies and equipment purchasing decisions. Conceptually, “network-planning models” refer to a collection of models that are used to determine how many passengers want to fly, which itineraries (defined as a flight or sequence of flights) they choose, and the revenue and cost implications of transporting passengers on their chosen flights.

Although various air carriers, aviation consulting firms and aircraft manufacturers own proprietary network-planning models, very few published studies exist describing them. Further, because the majority of academic researchers did not have access to the detailed ticketing and itinerary data used by airlines, the majority of published models are based on stated preference surveys and/or a high level of geographic aggregation. These studies provide limited insights into the range of scheduling decisions that network-planning models must support. Recent work by Coldren and Koppelman provides some of the first details into network-planning models used in practice (Coldren 2005; Coldren and Koppelman 2005a, b; Coldren et al. 2003; Koppelman et al. 2008).

2.4.2 Overview of Major Components of Network-Planning Models

As shown in Fig. 2.2, “network-planning models” refer to a collection of sub-models. First, an *itinerary generation algorithm* is used to build itineraries between each airport pair using leg-based air carrier schedule data obtained from a

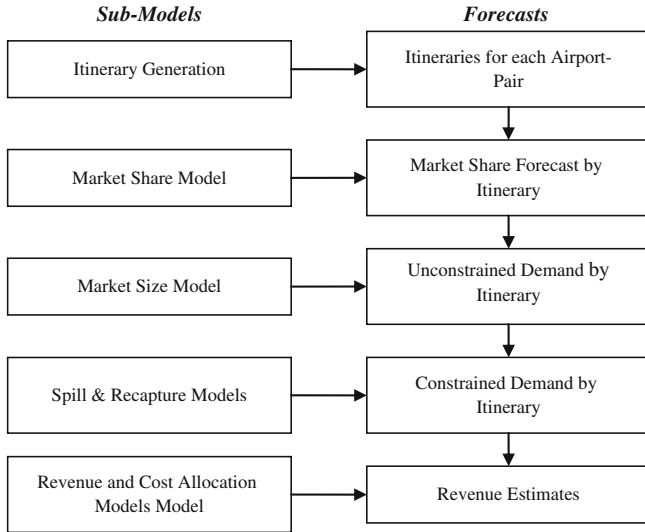


Fig. 2.2 Components and associated forecasts of a network-planning model

source such as the Official Airline Guide (OAG Worldwide Limited 2010). OAG data contain information for each flight including the operating airline, marketing airline (if a code-share leg), origin, destination, flight number, departure and arrival times, equipment, days of operation, leg mileage and flight time. Itineraries, defined as a flight or sequence of flights used to travel between the airport pair, are constructed from the OAG schedule. Itineraries are usually limited to those with a level-of-service that is either a non-stop, direct (a connecting itinerary not involving an airplane change), single-connect (a connecting itinerary with an airplane change) or double-connect (an itinerary with two connections). For a given day, an airport pair may be served by hundreds of itineraries, each of which offers passengers a potential way to travel between the airports. Although the logic used to build itineraries differs across airlines, in general itinerary generation algorithms include several common characteristics. These include distance-based circuitry logic to eliminate unreasonable itineraries and minimum and maximum connection times to ensure that unrealistic connections are not allowed. In addition, itineraries are typically generated for each day of the week to account for day-of-week differences in service offered.

An exception to the itinerary generation algorithm described above was developed by Boeing Commercial Airplanes for large-scale applications used to allocate weekly demand on a world-wide airline network. In this application, a weekly airline schedule involves the generation of 4.8 million paths across 280,000 markets that are served by approximately 950 airlines with 800,000 flights. Boeing’s algorithms, outlined in Parker et al. (2005), integrate discrete choice theory into both the itinerary generation and itinerary selection. That is, the

utility value of paths is explicitly considered as the paths are being generated; those paths with utility values “substantially lower” than the best path in a market are excluded from consideration.

After the set of itineraries connecting an airport pair is generated, a *market share model* is used to predict the percentage of travelers who select each itinerary in an airport pair. Different types of market share models are used in practice and can be generally characterized based on whether the underlying methodology uses a QSI or discrete choice (or logit-based) framework. Both types of market share models are discussed in this chapter.

Next, demand on each itinerary is determined by multiplying the percentage of travelers expected to travel on each itinerary by the forecasted *market size*, or the number of passengers traveling between an airport pair. However, because the demand for certain flights may exceed the available capacity, *spill and recapture models* are used to reallocate passengers from full flights to flights that have not exceeded capacity. Finally, *revenue and cost allocation models* are used to determine the profitability of an entire schedule (or a specific flight).

Market size and market share information can be obtained from ticketing data that provide information on the number of tickets sold across multiple carriers. In the U.S., ticketing data are collected as part of the U.S. Department of Transportation (U.S. DOT) *Origin and Destination Data Bank 1A or Data Bank 1B* (commonly referred to as DB1A or DB1B). The data are based on a 10% sample of flown tickets collected from passengers as they board aircraft operated by U.S. airlines. The data provide demand information on the number of passengers transported between origin–destination pairs, itinerary information (marketing carrier, operating carrier, class of service, etc.) and price information (quarterly fare charged by each airline for an origin–destination pair that is averaged across all classes of service). Although the raw DB datasets are commonly used in academic publications (after going through some cleaning to remove frequent flyer fares, travel by airline employees and crew, etc.), airlines generally purchase “Superset” data from the company Data Base Products (Data Base Products Inc. 2010). Superset data are a cleaned version of the DB data that are cross-validated against other data sources to provide a more accurate estimate of market sizes. See the websites of the Bureau of Transportation Statistics or Data Base Products for additional information.

The U.S. is the only country that requires airlines to report a 10% sample of used tickets. Thus, although ticketing information about domestic U.S. markets is publicly available, the same is not true for other markets. Two other sources of ticketing information include the Airline Reporting Corporation (ARC) and the Billing Settlement Plan (BSP), the latter of which is affiliated with the International Air Transport Association (IATA). ARC is the ticketing clearinghouse for many airlines in the U.S. and essentially keeps track of purchases, refunds and exchanges for participating airlines and travel agencies. Similarly, BSP is the primary ticketing clearinghouse for airlines and travel agencies outside the U.S.

Given an understanding of the major components of network-planning models and the OAG schedule, itinerary and ticketing data sources that are required to support the development of these models, the next sections provide a detailed description of QSI and logit-based market share models.

2.4.3 QSI Models

Market share models are used to estimate the probability a traveler selects a specific itinerary connecting an airport pair. Itineraries are the products that are ultimately purchased by passengers, and hence it is the characteristics of these itineraries that influence demand. In making their itinerary choices, travelers make tradeoffs among the characteristics that define each itinerary (e.g. departure time, equipment type(s), number of stops, route, carrier). Modeling these itinerary-level tradeoffs is essential to truly understand air-travel demand and is, therefore, one of the most important components of network-planning models.

The earliest market share models employed a demand allocation methodology referred to as QSI.¹ QSI models, developed by the U.S. government in 1957 in the era of airline regulation (Civil Aeronautics Board 1970) relate an itinerary's passenger share to its "quality" (and the quality of all other itineraries in its airport pair), where quality is defined as a function of various itinerary service attributes and their corresponding preference weights. For a given QSI model, these preference weights are obtained using statistical techniques and/or analyst intuition. Once the preference weights are obtained, the final QSI for a given itinerary is usually expressed as a linear or multiplicative function of its service characteristics and preference weights. For example, suppose a given QSI model measures itinerary quality along four service characteristics (e.g. number of stops, fare, carrier, equipment type) represented by independent variables X_1 , X_2 , X_3 , X_4 and their corresponding preference weights β_1 , β_2 , β_3 , β_4 . The QSI for itinerary i , QSI_i , can be expressed as:

$$QSI_i = (\beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_4 X_4), \text{ or}$$

$$QSI_i = (\beta_1 X_1)(\beta_2 X_2)(\beta_3 X_3)(\beta_4 X_4).$$

Other functional forms for the calculation of QSI's are also possible. For itinerary i , its passenger share is then determined by:

$$S_i = \frac{QSI_i}{\sum_{j \in J} QSI_j}$$

¹ QSI models described in this section are based on information in the Transportation Research Board's Transportation Research E-Circular E-C040 (Transportation Research Board 2002) and on the personal experiences of Gregory Coldren and Tim Jacobs.

where

S_i is the passenger share assigned to itinerary i ,

QSI_i is the quality of service index for itinerary i ,

$\sum_{j \in J} QSI_j$ is the summation over all itineraries in the airport pair.

Theoretically, QSI models are problematic for two reasons. First, a distinguishing characteristic of these models is that their preference weights (or sometimes subsets of these weights) are usually obtained independently from the other preference weights in the model. Thus, QSI models do not capture interactions existing among itinerary service characteristics (e.g. elapsed itinerary trip time and equipment, elapsed itinerary trip time and number of stops). Second, QSI models are not able to measure the underlying competitive dynamic that may exist among air travel itineraries. This second inadequacy in QSI models can be seen by examining the cross-elasticity equation for the change in the passenger share of itinerary j due to changes in the QSI of itinerary i :

$$\eta_{QSI_i}^{S_j} = \frac{\partial S_j}{\partial QSI_i} \frac{QSI_i}{S_j} = -S_i QSI_i.$$

The expression on the right side of the equation is not a function of j . That is, changing the QSI (quality) of itinerary i will affect the passenger share of all other itineraries in its airport pair in the same proportion. This is not realistic since, for example, if a given itinerary (linking a given airport pair) that departs in the morning improves in quality, it is likely to attract more passengers away from the other morning itineraries than the afternoon or evening itineraries.

Thus, to summarize, because QSI models have a limited ability to capture the interactions between itinerary service characteristics or the underlying competitive dynamic among itineraries, other methodologies, such as those based on discrete choice models have emerged in the industry. A detailed overview of discrete choice models is provided in the Customer Modeling chapter. An overview of how discrete choice models have been applied to market share modeling is presented in the next section.

2.4.4 Application of Discrete Choice Models to Market Share Modeling

As presented in the Customer Modeling chapter, discrete choice (or “logit”) models such as the multinomial logit (MNL) model are random utility maximizing models that describe how individuals choose one alternative among a finite set of mutually exclusive and collectively exhaustive alternatives. The individual chooses the alternative that has the maximum utility. The utility function for a random utility model is defined as

$$U_{ni} = \beta^* x_{ni}$$

where

U_{ni} is the total utility of alternative i for individual n .

β^* is the vector of parameters associated with attributes x . Utility is assumed to be a linear in parameters function of attributes x .

x_{ni} is the vector of attributes that vary across individuals and alternatives.

Because the utility the individual receives from each alternative is not known to the researcher, the utility function is assumed to have two components. The systematic or representative component contains observed variables that describe characteristics of the individuals and alternatives. The unobserved or error component is a random term that represents the unknown (to the researcher) portion of the individuals' utility function. The utility function is estimated using

$$V_{ni} = \beta x_{ni} + \varepsilon_{ni}$$

where

V_{ni} is the total observed utility of alternative i for individual n

β is the vector of estimates for β^*

x_{ni} is the vector of attributes for alternative i and individual n

ε_{ni} is an unobserved error component.

Different choice models are derived by imposing assumptions about the distribution of the error term and/or β . For example, the assumption that the error term is independent and identically distributed Gumbel² with mode³ 0 and scale γ , iid $G(0, \gamma)$ leads to the multinomial logit (MNL) model (McFadden 1974). The MNL probability of selecting alternative i among all j alternatives in C_n , the choice set for individual n , can be expressed in closed-form as

$$P_{ni} = P(i|x_{ni}, \beta) = \frac{e^{\beta' x_{ni}}}{\sum_j e^{\beta' x_{nj}}}.$$

The main limitation of the MNL is exhibited in the independence of irrelevant alternatives (IIA) property which states that the ratio of choice probabilities P_{ni}/P_{nj} for $i, j \in C_n$ is independent of the attributes of any other alternative. The IIA property of the MNL model is also apparent by examining the cross-elasticity equation for the change in the probability of itinerary j due to changes in an attribute of itinerary i :

² An iid Gumbel distribution is also called Type I extreme value.

³ Several publications incorrectly report the parameters describing the Gumbel distribution as the mean and scale. However, the shape and dispersion of the Gumbel distribution are formally defined by the mode and scale. Further, the mean of the Gumbel is given by the relationship $\text{mean} = \text{mode} + \{0.577/\gamma\}$.

$$\eta_{X_{ik}}^{P_j} = \frac{\partial P_j}{\partial X_{ik}} \frac{X_{ik}}{P_j} = -P_i \beta_k X_{ik}$$

Note that the expression on the right side is not a function of j . The IIA property of MNL is equivalent to the elasticity problem of the QSI model; that is, the cross-elasticity is undifferentiated across alternatives. In terms of substitution patterns, this means a change or improvement in the utility of one alternative will draw share proportionately from all other alternatives. In many applications, this may not be a realistic assumption. For example, in itinerary choice, the *unobserved factors* associated with the non-stop alternatives are expected to be correlated (since all non-stops are more convenient for passengers, may exhibit a decreased likelihood of lost luggage, etc.). Thus, the substitution between these alternatives is likely to be greater than between any of them and the connecting alternatives.

While the MNL model can be criticized for the restrictive substitution patterns it imposes, recent comparisons of itinerary choice models based on the MNL and QSI methodologies at a major U.S. airline clearly showed that the MNL outperformed their QSI model. In addition, many other discrete choice models (some developed specifically within the context of airline itinerary choice) can be used to incorporate flexible substitution patterns. Thus, the IIA property should not be viewed as a limitation of discrete choice models, as many other models (discussed extensively in Garrow 2010) can be used to relax this property. Nonetheless, it is interesting to note that in the context of itinerary choice models, even the simple MNL model dramatically outperformed the QSI model.

2.4.5 MNL and QSI Model Development at a Major U.S. Carrier

One of the first published studies modeling air-travel itinerary share choice based on a discrete framework was published in Coldren et al. (2003). MNL model parameters were estimated from a single month of itineraries (January 2000) and validated on monthly flight departures in 1999 in addition to selected months in 2001 and 2002. Using market sizes from the quarterly Superset data adjusted by a monthly seasonality factor, validation was undertaken at the flight segment level for the carrier's segments. That is, the total number of forecasted passengers on each segment was obtained by summing passengers on each itinerary using the flight segment. These forecasts were compared to onboard passenger count data. Errors, defined as the mean absolute percentage deviation, were averaged across segments for regional entities and compared to predictions from the original QSI model. Regional entities are defined by time zone for each pair of continental time zones in the U.S. (e.g. East–East, East–Central, East–Mountain, East–West,..., West–West) in addition to one model for the Continental U.S. to Alaska/Hawaii and one model for Alaska/Hawaii to the Continental U.S. The MNL forecasts were consistently superior to the QSI model, with the magnitude of errors reduced on the order of 10–15% of the QSI errors. Further, forecasts were stable across

months, including months that occurred after September 11, 2001. Additional validation details and estimation results are provided in Coldren et al. (2003).

The MNL itinerary share model reported in Coldren et al. (2003)—which represents a discrete choice model that was used to replace a major U.S. airlines' QSI model—includes variables for level-of-service, carrier presence, connection quality, aircraft size and type, fare and departure time-of-day. The representation of passengers' preferences for non-stop flights merits further discussion as it is unique from many other specifications found in practice and, more importantly, was found to be robust. Specifically, level-of-service (non-stop, direct, single-connect, or double-connect) is modeled using dummy variables to represent the level-of-service of the itinerary with *respect to the best level of service* available in the airport pair. That is, an itinerary with a double connection is much more onerous to passengers when the best level-of-service in the market is a non-stop than when the best level-of-service in the market is a single connection. Further, parameter estimates across 18 regional entities reveal that passengers have similar, but not identical responses, to changes in level-of-service across the entire domestic system. This is one example of the benefits of a “well-defined” utility function that captures the fundamental trade-offs of how passengers make itinerary choices; that is, parameter estimates are stable across datasets. Subsequently, this aids in transferability across different time periods and leads to better and more stable forecasting accuracy.

Another common industry practice reflected in the MNL itinerary share model of the major U.S. carrier is the inclusion of carrier presence variables. Numerous studies have found that increased carrier presence in a market leads to increased market share for that carrier (Algers and Beser 2001; Nako 1992; Prousaloglou and Koppelman 1999; Suzuki et al. 2001). In the MNL model, a “point of sale weighted airport presence” variable, used to represent carrier presence at both the origin and destination, is found to influence the value of itineraries.

Finally, it is important to note that in the major carrier's MNL itinerary share model, preferences for departure times are represented via the inclusion of time-of-day dummy variables for each hour of the day. In practice, there are other methods based on schedule delay formulations that are currently in use or are being explored in a research context to represent individuals' time-of-day preferences. Unfortunately, the terminology that has been used to describe the schedule delay functions is often referred to as a “nested logit model” within the airline community, which is incorrect. The next section clarifies the distinction between time-of-day preference and schedule delay functions. The Customer Choice chapter clarifies the definition of the nested logit model, as derived from discrete choice theory and discusses more advanced discrete choice models that have been applied to itinerary share and other airline applications.




Fig. 2.3 Passenger departure time preferences from time period model. *Source* Reprinted from Koppelman et al. (2008) with permission of Elsevier

2.4.6 Time-of-Day Preferences Versus Schedule Delay Functions

Determining when to schedule flights is arguably one of the most important decisions made by airline scheduling departments. Scheduling flights during unpopular departure times will result in fewer passengers and/or lower average fares. As described in Koppelman et al. (2008), different approaches have been used to model air travelers' departure time preferences. The first is to include time-of-day dummy variables for each hour of the day. Figure 2.3 shows an example of time-of-day preferences for a month of continental U.S. departures based on all flights traveling westbound by one time zone. Parameter estimates based on a MNL formulation indicate passengers prefer to depart early in the morning (at 8 a.m.) or later in the afternoon (5 p.m.).

However, the use of discontinuous time periods poses interpretation problems in practice as a slight change in schedule (e.g. from 9:59 p.m. to 10:02 p.m.) can cause large and counter-intuitive shifts in the probabilities. Estimation of parameters for all time periods can also be subject to over-fitting problems. To address these deficiencies, Koppelman et al. (2008) propose an approach adopted by Zeid et al. (2006) in the context of urban travel activity models. The approach is to estimate weighting parameters for a series of sine and cosine curves to obtain an overall representation of the distribution of departure time preferences. The time-of-day preference for three sine and cosine curves is specified as:




Fig. 2.4 Time of day utility curve from base sin-cos model. *Source* Reprinted from Koppelman et al. (2008) with permission of Elsevier

$$\begin{aligned} \text{Utility for alternative } i = & \beta_1 \sin\left(\frac{2\pi t_i}{1440}\right) + \beta_2 \sin\left(\frac{4\pi t_i}{1440}\right) + \beta_3 \sin\left(\frac{6\pi t_i}{1440}\right) \\ & + \beta_4 \cos\left(\frac{2\pi t_i}{1440}\right) + \beta_5 \cos\left(\frac{4\pi t_i}{1440}\right) + \beta_6 \cos\left(\frac{6\pi t_i}{1440}\right) \end{aligned}$$

where

t_i is the departure time of itinerary i expressed as minutes past midnight
1440 is the number of minutes in the day.

The final time-of-day value from this model is obtained by summing the six weighted trigonometric functions and is shown in Fig. 2.4. Statistical tests indicate that continuous specification is preferred over the time-of-day dummy variables.⁴

Carrier (2008) proposed a modification to this formulation to account for cycle lengths that are shorter than 24 hours. Formally, the equation:

$$\beta_1 \sin\left(\frac{2\pi h}{1440}\right) + \beta_2 \cos\left(\frac{2\pi h}{1440}\right) + \dots$$

is replaced with

⁴ Other trigonometric functions involving an additional offset parameter, such as those proposed by Gramming et al. (2005) were also estimated as part of the analysis. Results from these two approaches were virtually identical.

$$\beta_1 \sin\left\{\frac{2\pi(h-s)}{d}\right\} + \beta_2 \cos\left\{\frac{2\pi(h-s)}{d}\right\} + \dots$$

$$l - e \leq d \leq 24$$

$$0 \leq s \leq e$$

where e and l represent the departure times of the earliest and latest itineraries in the market, respectively, h represents the departure time, s represents the start time of the cycle (which is not uniquely identified and can be set to an arbitrary value) and d represents the cycle duration. The examples in this chapter use the 24-hours period, as Carrier's formulation leads to a nonlinear-in-parameters function, which he solved using a trial-and-error method. The trial-and-error method (often used by discrete choice modelers when they encounter nonlinear-in-parameters functions) essentially fixes d to different values and estimates the remaining parameters. The value of d that results in the best log likelihood value is the preferred model.

Figures 2.3 and 2.4 represent time-of-day preferences on a 24-hours cycle by measuring the relative value of a departure time relative to all other possible departure times. However, from a behavioral perspective, itinerary selection may be influenced by an individual's effort to depart as close as possible to his/her ideal departure time. The difference between an individual's desired departure time and actual departure time is defined as schedule delay. Formally, schedule delay for itinerary i , SD_i , can be expressed as:

$$SD_i = \sum_{\forall j} g(DT_i - TP_j) W_{TP_j}$$

where

DT_i is the departure time for itinerary i

TP_j is the start of each 15-min time period from 5:30 a.m. to 10:30 p.m.

$g()$ is a transformation function of the difference in minutes between the itinerary departure time and the time period

W_{TP_j} is the weight for time period j .

There are two key points to note about this formulation. First, the weights W_{TP_j} account for departure preferences and the distribution of observed passenger departures at different times-of-day. Second, the formulation is general in the sense that different schedule delay transformation functions are possible. Several functions, including linear, square root, squared, logistic, etc. were estimated. The logistic transformation shown in Fig. 2.5 was found to fit the data the best. Formally, this schedule delay transformation and time period weights are given as:

$$g(DT_i - TP_j) = \frac{1}{1 + \exp\left(\frac{\alpha_2 - |DT_i - TP_j|}{\alpha_1}\right)}$$

$$W_{TP_j} = \frac{\text{sincosValue}_{TP_j}}{\sum_j \text{sincosValue}_{TP_j}}$$




Fig. 2.5 Impact of logit function schedule delay on itinerary value *Source:* Reprinted from Koppelman et al. (2008) with permission of Elsevier

where

DT_i is the departure time for itinerary i

TP_j is the start of each 15-min time period from 5:30 a.m. to 10:30 p.m.

$\text{sincosValue}_{TP_j}$ is the sum of the added terms for time period j .

The formulation based on schedule delay is found to fit the data better than the model based on continuous time-of-day preferences. Additional results that capture differences in day-of-week departure time preferences as well as early and late departure (or arrival) delays by outbound and inbound itineraries are reported in Koppelman et al. (2008).

2.4.7 Summary

This section focused on describing two major types of market share models found in scheduling models: those based on the QSI methodology and those based on discrete choice methods. An emphasis was placed on identifying concepts that in the authors' experiences are commonly misunderstood in practice. This includes the definition of nested logit models and schedule delay functions.

Based on our interviews with industry experts, we learned that many airlines currently using logit-based methodologies are contemplating reintroducing QSI methodologies due to the perceived complexity of logit models and difficulty in maintaining parameter estimates. However, in our opinion, we believe that the fundamental problems currently being observed are not due to the use of a logit model, but rather over-parameterized utility functions. One of the primary differences between the published MNL model of a major U.S. carrier (which was

clearly seen to dominate their QSI model) and the logit models used in practice relates to the number of variables (and estimated β coefficients).

In the published MNL model, each regional entity has 36 parameter estimates in addition to estimates associated with each airline carrier. Further, 18 of these parameters, which are associated with dummy variables for time-of-day preferences, can be further reduced via incorporation of an appropriate continuous schedule delay function. This is in comparison with alternative logit models reported to have hundreds, *if not thousands*, of parameter estimates. However, a simple, yet well-specified MNL utility function can lead to superior predictive performance over a QSI model. Complexity should not be driven by the number of variables included in the model, but rather by the desire or need to obtain more accurate substitution patterns than those imposed by the MNL. Further, more flexible patterns can be incorporated via the use of more advanced GEV or mixed logit models discussed in the Customer Modeling chapter.

2.5 Airline Fleet Assignment Process and Schedule Development

2.5.1 Introduction and Scope

The fleet assignment process represents one of the most important and well studied applications of operations research in the airline industry. In many ways the schedule development and fleet assignment process embodies the complexities and computational difficulties characteristic of many aspects of the airline industry. To begin, many carriers use the fleet assignment process to help finalize market frequencies, flight times and enforce various operational requirements on the schedule. These may include operational needs such as station purity in which particular stations are limited to one or two types of fleet to meet maintenance and engineering capabilities, the incorporation of minimum revenue guarantees (MRG) in which municipalities contract for service to their airport, and the increase or reduction of available aircraft due to retirements and new deliveries.

Later in the schedule development process, the fleet assignment process and optimization tools are used to finalize the fleet assignments, distribute various sub-fleets within the network based on operational limitations such as range, and incorporate maintenance opportunities and crew considerations. For example, a carrier might fly several markets with a Boeing 737 but some of the markets may require a 737-800 rather than a 737-200 due to range limitations. Incorporating maintenance opportunities may involve having a specified number of aircraft of a specific type on the ground for 12 hours beginning between 1800 (6:00 p.m.) and 2000 (8:00 p.m.) in the evening to ensure that enough aircraft are available to launch operations the following morning. The carrier may also want their flight crews to stay with the same aircraft for as long as possible to minimize “crew connections” in which the crew leaves one aircraft upon arrival and flies another

aircraft for their next scheduled flight. Having the crew stay with the aircraft saves time for both the crew and the airline and results in a more efficient operation and better utilization of the aircraft. It also facilitates a more effective line maintenance operation during the operating day due to the opportunity for maintenance personnel to discuss issues with the crew during aircraft turns when needed.

The efficient utilization of expensive resources is an objective of any profitable airline. One important aspect of this utilization process is fleet assignment. Fleet assignment involves the optimal allocation of a limited number of fleet types to flight legs in the schedule subject to various operational constraints. The most common form of the FAM makes simplifying assumptions about passenger demands, revenues and network flows to approximate the expected revenue for each flight leg in the schedule. These simplifying assumptions provide a point estimate of the expected revenue for each leg in the schedule given various capacity options. The following section presents the basic development of the most common form of the fleet assignment model. In addition, the following section will also present two potential enhancements to the typical fleet assignment model that incorporates the O&D passenger flows into the process. Following the development of the fleet assignment model and its enhancements, we compare and contrast two formulations and present example results using actual airline schedules.

2.5.2 Fleet Assignment Model Development

The fleet assignment problem is typically posed as a binary assignment model in which a specific aircraft type is assigned to each leg in the schedule. The basic fleet assignment model maximizes overall profit subject to three primary constraints: (1) plane count, (2) balance and (3) cover. The plane count constraint stipulates that, for each fleet type, the number of planes used to fly the schedule cannot exceed the total number of planes available. The balance constraint stipulates that the number of arrivals must equal the number of departures for each station, time event and fleet type. The cover constraint requires that a fleet type be assigned to each leg in the network. Most airlines include numerous additional operational constraints that help tailor the solution to the specific operational requirements of the airline.

Typically, fleet assignment models pose the problem as a multi-commodity flow problem. For the fleet assignment problem, arcs represent the arrival and departure of flights and aircraft on the ground. The nodes define the specific time and station where these activities take place. Figure 2.6 illustrates the basic timeline approach used by the fleet assignment model. Figure 2.6a portrays the actual timeline of flights arriving and departing a single station. Figure 2.6b presents the node/arc representation of the timeline for the same station. Figure 2.6c provides a detailed schematic of the decision variables that represent the selection of aircraft type i assigned to flight leg j and ground arcs flowing into a single node at time t and station s within the timeline.

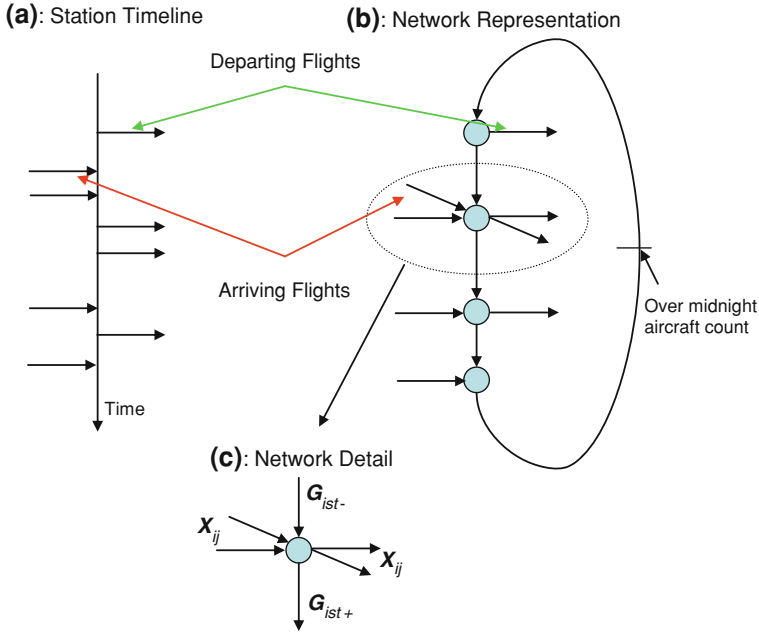


Fig. 2.6 Network representation of a typical FAM formulation. **a** Station timeline. **b** Network representation. **c** Network detail

Using this network representation, we formally develop the basic FAM proposed by Hane et al. (1995) using notation similar to that used by Lohatepanont (2001) and Smith and Johnson (2006).

2.5.2.1 Definition of Sets

S : set of stations or airports indexed by s .

J : set of flight legs indexed by j .

F : set of fleet types (e.g. S80, 737) indexed by i .

T : set of all departure and arrival events indexed by t .

$Re(i)$: set of all flight legs for fleet type i crossing the counting line (e.g. midnight) indexed by j .

$IN(i, s, t)$: set of flight legs inbound to $\{i, s, t\}$.

$OUT(i, s, t)$: set of flight legs outbound from $\{i, s, t\}$.

2.5.2.2 Decision Variables

$$x_{ij} = \begin{cases} 1 & \text{if aircraft type } i \in F \text{ is assigned to schedule leg } j \in J, \\ 0 & \text{otherwise} \end{cases}.$$

G_{ist^+} represents the number of aircraft on the ground for fleet type $i \in F$, at stations $s \in S$, on the ground arc just following time $t \in T$.

G_{ist^-} represents the number of aircraft on the ground for fleet type $i \in F$, at stations $s \in S$, on the ground arc just prior to time $t \in T$.

2.5.2.3 Parameter Definitions

R_{ij} represents the expected revenue associated with assigning aircraft type $i \in F$ to flight leg $j \in J$ and is a function of expected demand, spill and unit revenue per passenger.

C_{ij} represents the expected costs associated with assigning aircraft type $i \in F$ to flight leg $j \in J$ as a function of fixed, ownership and variable costs.

NP_i represents the number of available aircraft of type $i \in F$.

2.5.2.4 Conventional Leg-Based FAM Formulation

$$\max P = \sum_{j \in J} \sum_{i \in F} (R_{ij} - C_{ij}) x_{ij} \quad (\text{Objective : Maximize Profit}) \quad (2.1)$$

subject to:

$$\sum_{j \in Re(i)} x_{ij} + \sum_{s \in S} G_{is0^-} \leq NP_i \quad \forall i \in F \quad (\text{Plane Count}) \quad (2.2)$$

$$G_{ist^-} - G_{ist^+} + \sum_{j \in IN(i,s,t)} x_{ij} - \sum_{j \in OUT(i,s,t)} x_{ij} = 0 \quad \forall i \in F, s \in S, t \in T \quad (\text{Balance}) \quad (2.3)$$

$$\sum_{i \in F} x_{ij} = 1 \quad \forall j \in J \quad (\text{Cover}) \quad (2.4)$$

$$\begin{aligned} x_{ij} &\in \{0, 1\} \quad \forall i \in F, \quad \forall j \in J \\ G_{isj} &\geq 0 \quad \forall i \in F, \quad s \in S, \quad t \in T \end{aligned} \quad (2.5)$$

Constraints (2.2) represent resource constraints and states that the number of planes of each fleet type i cannot exceed the total number of planes available, NP_i . Constraints (2.3) represent the balance constraints stating that, at any station and time, the arrival of an aircraft must be matched by the departure of the aircraft. Aircraft can arrive at a station from another station or from the same station in the previous time event, $t - 1$. The time events in this formulation represent an arrival or departure event or a combination of arrivals and departures at the station. Constraints (2.4) represent the cover constraints that stipulate each flight leg must be assigned a fleet type. Constraints (2.5) define the decision variable for assigning fleet type i to flight leg j as a binary variable and specify non-negativity for the ground arcs. Several references present overviews of the general leg-based FAM (see Abara 1989; Subramanian et al. 1994; Hane et al. 1995).

The conventional fleet assignment model described by Eqs. 2.1–2.5 is used for both long-term planning of the airline schedule and near-term finalization of the schedule fleet allocation. Depending on the carrier, this model can be used for long-term planning to fleet a typical daily schedule or a weekly schedule. Most U.S. based carriers tend to plan schedules based on a typical day while European and Asian carriers tend to focus on weekly schedules.

During the planning process, carriers will often need to reduce or expand their schedules to better match available capacity to demand during off seasons such as the fall and early winter or high demand seasons such as Christmas and New Year's holidays. In these cases, airlines can use the FAM to help select and fleet flights that best contribute to the overall profitability of the schedule while dropping other flights that do not contribute. To reduce the schedule, airlines can run the FAM in “reduction mode” in which we relax the cover constraint described by Eq. 2.4. Relaxing Eq. 2.4 as a less than or equal to constraint allows the model to drop flights that do not help maximize the overall profitability of the schedule.

Similarly, the relaxed version of the FAM can be used to expand the schedule to capture the need for more capacity during high demand seasons. To accomplish this, many airlines “overbuild” schedules to include more flights than they expect to operate. Using the overbuilt schedule as input, the airline then optimizes the schedule using the FAM in reduction mode to drop less desirable flights from the schedule.

In both cases mentioned above, the airline must add a number of operational constraints to the FAM to prevent undesirable results. Often these added constraints place bounds on the number of frequencies allowed in any market and limit the number of flights dropped overall. This prevents the model from dropping out of markets entirely or radically reducing the resulting schedule in the name of expected profitability. One caveat that should be kept in mind involves the revenues and costs used to drive the objective function when using reduction mode. The FAM described by Eqs. 2.1–2.5 requires accurate reflections of the revenues and costs associated with each potential fleet/flight combination considered by the

model. These revenue and cost estimates are dependent on the original schedule used to forecast demand and passenger traffic. Relaxing the cover constraint (Eq. 2.4) compromises the accuracy of these forecasts and the overall estimates for revenue and costs. Many carriers try to limit this undesirable impact by iteratively updating the forecasted revenues and costs to ensure more accurate results. However, this problem extends beyond schedule reductions. In fact, the revenues, costs and fleetings results can be influenced by the overall mix of local and connecting passengers. As a result, several enhancements to the FAM have been proposed to incorporate the influence of connecting passengers and accurately reflect the change in revenues due to limited schedule reductions.

During the near-term planning process, the FAM can be used to finalize the overall fleetings (allocate sub-fleets), incorporate crew considerations into the final schedule and build transition schedules that bridge one seasonal schedule into the next. To build a transition schedule, an airline can formulate the FAM using the final fleet assignments from the two seasonal schedules as inputs and allow the model to optimize the fleet assignments to connect the two schedules. In addition, the FAM can be used to re-fleet portions of the schedule to better match overall demand to available capacity near the day of departure. We present an actual case study of this type of application near the end of the chapter.

As highlighted above, the major problem with the leg-FAM approach described by Eqs. 2.1–2.5 is that it does not accurately incorporate the O&D marketing effects and expected passenger flows throughout the network. The fleet assignment process should account for multiple markets utilizing each leg of the schedule, multiple classes within each market, and network interactions caused by the various markets competing for space.

Several approaches to incorporating RM aspects into FAM have been investigated over the past 10 years to develop an Origin–Destination Passenger-based Fleet Assignment Model (ODFAM). These approaches have dealt with the size and non linearity of ODFAM through various decomposition approaches.

Farkas (1996) demonstrates that RM has a significant impact on traffic volume and mix and by ignoring these effects FAM can yield sub-optimal solutions. His analysis illustrated the necessity of modeling the effects of both network flow and stochastic demand to improve FAM performance. He concludes that incorporating RM directly into FAM is not practical. He proposes three approaches to this problem:

- Column generation. Where each column represents a complete fleetings solution. The master evaluates traffic and revenue, and ensures that allocations do not exceed capacity. The columns are generated using a multi-commodity formulation. Although no computational results are published Farkas states that the subproblem is relatively slow to solve (40 min) and is impractical for operational use.
- Leg Class revenue management FAM. Since many airlines do not have full network control in their RM systems, Farkas investigates the impact of leg class revenue management control on FAM. He shows that for a typical airline fare

structure, the revenue function could be non-concave. This non-concavity makes this formulation unattractive in terms of computational efficiency.

- Decomposing the flight schedule into subnetworks between which there are limited or no leg-interactions. Fleeting solutions for each sub-network are generated, the traffic and revenue for each sub-network is evaluated with a Monte Carlo simulation. In the FAM formulation, each of the assignments for a subnetwork is represented by one meta-variable. By starting with a feasible leg-FAM solution, this approach should always produce improving solutions. No computational results are available.

Knicker (1998) and Barnhart et al. (2002) investigate the interactions between RM and FAM. In this work, the authors develop a Passenger Mix Model (PMM) that gives a schedule with known flight capacities and a set of passenger demands with known fare, and determines optimal traffic and revenue. PMM includes aspects of customer choice modeling and includes recapture (the probability that a customer who is spilled from one flight leg books on another of the same airline). PMM assumes that demand is deterministic and that the airline has complete knowledge and control of which passengers they accept. PMM could be formulated as a multi-commodity flow problem but due to the large number of passenger types and potential paths this approach is impractical. Knicker reduces the problem by using key-paths, the originally desired itinerary for each passenger. Alternate itineraries are necessary only when passengers are spilled from their preferred itinerary. The problem is solved using column generation, with each column representing passengers spilled from one itinerary and recaptured on another. Knicker formulates the stochastic version but does not present results.

Knicker combines PMM and FAM. The integrated problem, IFAM, is solvable but suffers from increased fractionality versus leg-FAM in which aggregate leg revenues and costs are used to reflect profitability of different fleets on a flight leg. He improves performance through coefficient reduction and additional cuts, but the MIP is still much more difficult to solve than the corresponding leg-FAM MIP. Knicker compares performance of various approaches using a Monte Carlo simulation model. By comparing models that capture the network effects assuming deterministic demand versus stochastic models that ignore network effects, he shows that if flow demand is at least 25% of the total demand, then capturing network effects is more important than capturing stochastic effects. Knicker does not formulate a version of FAM that addresses both stochastic demand and network effects.

Lohatepanont (2001) continues the analysis of IFAM. He investigates the sensitivity of IFAM to several of the simplifying assumptions in its formulation:

- *Demand uncertainty.* IFAM assumes that demand is fixed and known. The demands used in FAM are forecasts subject to random and systematic errors
- *Imperfect control.* PMM assumes that airlines have complete control over which passengers are accommodated
- *Recapture rate errors.* PMM assumes that recapture rate is known.

Table 2.1 Itinerary demand information

Market	Itinerary (sequence of flights)	Number of passengers	Average fare
ABQ-DFW	001	75	\$220
DFW-BOS	002	120	\$225
ABQ-BOS	001–002	80	\$400

Through simulation analysis of IFAM and PMM Lohatepanont shows that while relaxing these assumptions, to make the models more realistic, reduces the benefit of IFAM versus leg-FAM, IFAM consistently outperforms FAM. Barnhart et al. (2002) provide an excellent recapitulation of Kniker and Lohatepanont’s work and the relationship between capacity assignments and RM passenger allocations in a deterministic setting. We present the model proposed by Barnhart et al. (2002) in the next section.

Erdmann et al. (1997) proposes a sequential approach to the itinerary FAM problem. They solve FAM and then the passenger mix problem. Klierer (2000) proposes an approach that integrates FAM and RM using simulated annealing. Klierer uses a neighborhood search strategy, starting with an initial feasible solution and looks for improving assignment swaps. He accepts or rejects new solutions based on a simulated annealing strategy. The revenue is evaluated with a deterministic passenger flow model.

The model proposed by Barnhart et al. 2002 improves the conventional leg-based FAM by explicitly incorporating the *network* and *recapture effects* into the fleeing process. To better understand this motivation, consider the following example.

Network and Recapture Effects: An Illustrative Example

Consider a small airline network with two flight legs: Flight 001 ABQ-DFW and Flight 002 DFW-BOS. Table 2.1 shows demand and fare data in three OD markets ABQ-DFW, DFW-BOS and ABQ-BOS (connecting through DFW). If 100-seat aircraft is assigned to both flight legs in the network, the optimal revenue for the network of \$41,875 is obtained by accommodating 75, 75 and 25 passengers from ABQ-DFW, DFW-BOS, and ABQ-BOS markets respectively.

In conventional FAM, flight legs are assumed to be independent; consequently, fares of connecting passengers have to be allocated to corresponding flight legs in the itineraries. Knicker (1998) experiments with a number of fare allocation schemes and shows that no single allocation scheme, which is applicable to all networks, exists. In this example, we use a simple “equal-fare” allocation, in which the connecting fare is divided equally among the flight legs making up the itinerary. Thus, in our example, the ABQ-BOS fare of \$400 is equally divided and allocated to ABQ-DFW and DFW-BOS flights (\$200 each). With this allocation,

the optimal leg-based revenue is obtained by maximizing the revenue on each flight independently. As a result, the optimal passenger mix is 75 and 25 passengers from ABQ-DFW and ABQ-BOS markets respectively for Flight 001, and 100 passengers from DFW-BOS market alone for Flight 002. The optimal revenue is \$44,000. Notice that the resulting optimal mix of passenger is infeasible because none of the ABQ-BOS passengers get on Flight 002, and thus the revenue of \$44,000 is inaccurate and unachievable.

Alternatively, one can view this as leg-based FAM's inability to calculate *spill* consistently in the network. When the total passenger demand for a flight leg exceeds the capacity of that flight leg, some passengers are not accommodated or are *spilled*. In this example, with leg-based FAM, 55 ABQ-BOS passengers are spilled from Flight 001, but 80 ABQ-BOS passengers are spilled from Flight 002. On the other hand, with the optimal passenger mix given at the beginning of this example, 55 ABQ-DFW passengers are consistently spilled from both legs.

Next, we introduce the concept of recapture. Normally, spilled passengers are either (1) lost to the airline (that is, they choose to travel on competing airlines or choose not to travel by air) or (2) *recaptured* on alternative flights in the network of the original airline. These recaptured passengers generate *recaptured revenue* for the airline on alternative flights. Most leg-based fleet assignment models ignore totally these recaptured revenues in their estimation of flight-leg revenues because inconsistent spills cannot possibly lead to accurate recapture estimates.

In conclusion, leg-based FAM cannot estimate flight-leg revenues accurately because it assumes flight-leg independency. Specifically, the inaccuracy is a result of (1) the inconsistent estimates of spills due to the *network effect* and, consequently, (2) the inaccurate estimates of (possibly significant) recapture revenues due to the *recapture effect*.

Further, this example demonstrates how decisions made independently for each flight leg are incorrect and suboptimal. To get an accurate estimate of passenger revenue, one needs to take a holistic look at the entire network.

Passenger Mix Model

As the example above shows, a leg-based view of the network does not accurately capture the network effects due to O&D passenger flows and recapture. We need a tool that can estimate the network revenue more accurately. Specifically, we need a tool that can estimate spill consistently throughout the network and allow recaptured revenue to be estimated systematically. Knicker (1998) proposes the PMM for this purpose. The objective of PMM is to find the optimal itinerary-based mix of passengers that maximizes the total revenue (including recaptured revenue) or, equivalently, minimizes the total *spill cost*, the revenue loss due to spilled passengers. PMM is formulated as follows:

$$\text{Min } \sum_{p \in P} \sum_{r \in P} \left(\text{fare}_p - b_p^r \text{fare}_r \right) \cdot t_p^r \quad (2.6)$$

$$\text{Subject to : } \sum_{p \in P} \sum_{r \in P} \delta_i^p t_p^r - \sum_{r \in P} \sum_{p \in P} \delta_i^p b_r^p t_r^p \geq Q_i - \text{CAP}_i \quad \forall i \in L \quad (2.7)$$

$$\sum_{r \in P} t_p^r \leq D_p \quad \forall p \in P \quad (2.8)$$

$$t_p^r \geq 0 \quad \forall p, r \in P. \quad (2.9)$$

This formulation utilizes a special set of variables (*keypath* variables t_p^r), first proposed by Barnhart et al. (1995), to enhance model solution. Specifically, the model a priori assigns passengers to their desired itineraries; next, if the capacities on some flights are insufficient, the model finds an optimal way to spill passengers off from these flights such that the total *spill cost* (the revenue loss due to spilled passengers) is minimized. This model incorporates recaptures using a set of Quantitative Service Index (QSI) based parameters called recapture rates, b_p^r , which is defined as the recapture rate from itinerary p to itinerary r or the fraction of passengers spilled from itinerary p that the airline succeeds in redirecting to itinerary r .

Let P be the set of all itineraries and L be the set of all flight legs. The decision variable, t_p^r , is the number of passengers who are redirected from their desired itinerary p to an alternative itinerary r . The parameter fare_p denotes the averaged fare for itinerary p . The objective function (Eq. 2.6) minimizes the total spill cost $\left(\sum_{p \in P} \sum_{r \in P} \text{fare}_p \cdot t_p^r \right)$ less the recaptured revenue $\left(\sum_{p \in P} \sum_{r \in P} b_p^r \text{fare}_r \cdot t_p^r \right)$. Constraints (2.7) ensure that the total number of passengers for each flight leg i (which equals to the original passengers desiring this flight, Q_i , less the total passengers spilled from this flight, $\sum_{p \in P} \sum_{r \in P} \delta_i^p t_p^r$, plus the total passengers recaptured from other flights, $\sum_{r \in P} \sum_{p \in P} \delta_i^p b_r^p t_r^p$) does not exceed the capacity of that flight, CAP_i . δ_i^p equals 1 if itinerary p utilizes flight leg i , 0 otherwise. Constraints (2.8) and (2.9) ensure the number of spilled passengers for each itinerary does not exceed the demand for that itinerary and is not less than zero.

PMM is a large-scale LP model that requires specialized solution algorithm. Knicker (1998) proposes a column and row generation-based algorithm for the model. Specifically, only a small set of variables is included in the original master problem and, as the algorithm progresses, more columns (variables redirecting passengers to alternative itineraries) are generated as necessary. Notice also that in the optimal solution most of Constraints (2.8) are not binding because most of the passengers are traveling on their desired (*keypath*) itineraries. Thus, Constraints (2.8) can be initially omitted and subsequently generated back in as necessary.

Itinerary-Based Fleet Assignment Model

The Itinerary-Based Fleet Assignment Model (IFAM) is the integration of the PMM and the leg-based FAM. The IFAM formulation is:

$$\text{Min } \sum_{k \in K} \sum_{i \in L} c_{k,i} f_{k,i} + \sum_{p \in P} \sum_{r \in P} \left(\text{fare}_p - b_p^r \text{fare}_r \right) \cdot t_p^r \quad (2.10)$$

$$\text{Subject to : } \sum_{k \in K} f_{k,i} = 1 \quad \forall i \in L \quad (2.11)$$

$$y_{k,o,t^-} + \sum_{i \in IN(k,o,t)} f_{k,i} - y_{k,o,t^+} + \sum_{i \in OUT(k,o,t)} f_{k,i} = 0 \quad \forall k, o, t \in N \quad (2.12)$$

$$\sum_{o \in A} y_{k,o,t_m} + \sum_{i \in CL(k)} f_{k,i} \leq N_k \quad \forall k \in K \quad (2.13)$$

$$\sum_{k \in K} \text{SEAT}_k \cdot f_{k,i} + \sum_{p \in P} \sum_{r \in P} \delta_i^p t_p^r - \sum_{r \in P} \sum_{p \in P} \delta_i^p b_r^p t_r^p \geq Q_i \quad \forall i \in L \quad (2.14)$$

$$\sum_{r \in P} t_p^r \leq D_p \quad \forall p \in P \quad (2.15)$$

$$f_{k,i} \in \{0, 1\} \quad \forall k \in K, \forall i \in L \quad (2.16)$$

$$y_{k,o,t} \geq 0 \quad \forall k, o, t \in N \quad (2.17)$$

$$t_p^r \geq 0 \quad \forall p, r \in P. \quad (2.18)$$

The objective function (Eq. 2.10) minimizes the total operating cost $\left(\sum_{k \in K} \sum_{i \in L} c_{k,i} f_{k,i} \right)$

and the total spill (less recapture) cost $\left(\sum_{p \in P} \sum_{r \in P} \left(\text{fare}_p - b_p^r \text{fare}_r \right) \cdot t_p^r \right)$. Constraints

(2.11)–(2.13) are original FAM constraints—coverage, balance and count constraints, respectively. Constraints (2.14) are capacity constraints, which dictate that for a given flight leg i the capacity of the chosen assignment $\left(\sum_{k \in K} \text{SEAT}_k \cdot f_{k,i} \right)$ must

exceed the total traffic $\left(Q_i - \sum_{p \in P} \sum_{r \in P} \delta_i^p t_p^r + \sum_{r \in P} \sum_{p \in P} \delta_i^p b_r^p t_r^p \right)$. Constraints (2.15)

guarantee that no spills exceed demands. Constraints (2.17) ensure binary selectivity. And finally Constraints (2.17)–(2.18) ensure non-negativity.

Because IFAM is an integration of two large-scale models, Barnhart et al. (2002) propose a column and row generation-based solution algorithm for solving IFAM. Specifically, the column and row generations are applied to the PMM part of the model, that is, to the traffic variables (t_p^r) and demand constraints (Constraints (2.8)).

Barnhart et al. (2002) test the model and algorithm on actual large-scale data set of a major U.S. airline. Their results indicate significant savings from the optimal assignment of aircraft types to flight legs, taking into account network and recapture effects. Further, they present experiments to validate IFAM's three key assumptions/parameters, namely, (1) deterministic demand, (2) recapture rate and (3) optimal control of passenger mix through PMM. Their findings indicate that (1) in a simulation test using stochastic demand generator, IFAM fleet decisions consistently outperform FAM fleet decisions, (2) IFAM fleet decisions are not particularly sensitive to a reasonable range of recapture rates and (3) in a simulation test where suboptimal control of passenger mix is simulated, IFAM fleet decisions again consistently outperform FAM decisions. For detailed information and discussion, readers are referred to Barnhart et al. (2002).

Another approach for incorporating both network effects and the stochastic nature of demand first proposed by Jacobs et al. (1999, 2000, 2008) uses revenue management controls to drive the fleet decision process. This model uses a Benders decomposition approach to integrate the FAM model with a stochastic O&D Revenue Management model. We refer to this approach as O&D FAM. The revenue associated with any FAM solution depends on the capacity assignment for all flight legs. Given an assignment solution, the O&D revenue is estimated by the O&D RM sub-problem. The revenue function for the entire network is approximated in the master problem (FAM) using a series of Benders cuts. Each cut improves the accuracy of the revenue approximation in the master FAM problem. Once a specified accuracy is achieved in the relaxed master problem, the assignment variables are changed to integer variables and the MIP is solved.

This approach is appealing because it addresses both passenger flows within the network and demand uncertainty. It also provides a method of incorporating the passenger mix optimization model used for revenue management directly into the fleet assignment process.

Typically, airlines estimate the expected revenue for each fleet-flight combination using a proportional spill model (Swan 1983) and an average fare per passenger. This process applies a spill model to the total demand for each leg in the schedule individually. As a result, the leg-based revenue and profit estimates cannot capture the effects of network flow on the traffic of individual legs. Using this formulation, the total revenue for each leg in the schedule reflects an independent point estimate of the revenue function. This leads to errors in estimating the expected traffic and revenue for each leg in the schedule.

In reality, the revenue function accounts for the cumulative effect of all market classes flowing over leg j as a function of its capacity and incorporates the interaction between all the legs in the schedule. The revenue function is actually a concave function with respect to leg capacity, CAP_j , resulting from O&D network flow. Figure 2.7 presents an illustrative example of the revenue function for a single leg in the network.

Network flow or O&D yield management (O&D YM) solutions yield a set of bidprices for each leg which represents the dual value of the O&D YM capacity constraint and equals the slope of the revenue function at a given leg capacity,




Fig. 2.7 Expected revenue as a function of leg capacity

CAP_j (Fig. 2.3). This slope can be used to define a linear approximation and upper bound to the revenue function. Mathematically, this upper bound is expressed as:

$$R_{0_j} + \lambda_j CAP_j \geq R(CAP_j) \quad (2.19)$$

where R_{0_j} represents the right-hand side of the linear approximation to the revenue function and $R(CAP_j)$ equals the total revenue as a function of the capacity of leg j , CAP_j . λ_j defines the marginal value of an extra seat on leg j (the bidprice) resulting from the O&D yield management. Mathematically, the bidprice for leg j is defined by:

$$\lambda_j = \frac{\partial R(CAP_j)}{\partial CAP_j} \quad (2.20)$$

In practical terms, the bidprice represents the minimum acceptable price of a seat. More importantly, the bidprice represents the change in the *total system revenue* due to a unit change in the capacity of leg j . Therefore, the bidprice captures the cumulative effects of the market classes flowing over leg j and the interactions between leg j and the other legs in the network. Please see Appendix A for a complete formulation and review of the O&D Yield Management (O&D YM) model formulation.

For any solution of FAM and a corresponding set of bidprices, the total revenue for the schedule is the sum of the revenues realized on each leg. Summing over all the legs in the network yields the following upper bound on the total revenue (R_{Total}):

$$\sum_{j \in J} R_{0_{jv}} + \sum_{j \in J} \lambda_{jv} CAP_{jv} \geq R_{\text{Total}} \quad \forall v \in V \quad (2.21)$$

where v is the index in the set V for a specific FAM and YM solution. This relationship represents a Bender's cut (Parker and Rardin 1988; Nemhauser and Wolsey 1988; Bradley et al. 1977) and defines an overall upper bound on the total revenue for the schedule as a function of network flow results. Therefore, the overall revenue used in FAM is limited by a function of O&D passenger flow resulting from the O&D YM process.

Constraints (2.21) relate the FAM to the network flow model. This relationship allows decomposition of the O&D FAM model into two separate but related problems: (1) a linear fleet assignment model and (2) a nonlinear network flow model. Separately, each of these models can be solved using conventional IP or NLP methods. Using Constraints (2.21), the general FAM is modified to include O&D effects. The resulting *linear* FAM used by O&D FAM is defined as:

Linear FAM Formulation

$$\max P = R_{\text{Total}} - C_{\text{Total}} \quad (\text{Objective : Maximize Profit}) \quad (2.22)$$

subject to:

$$\sum_{j \in \text{Re}(i)} x_{ij} + \sum_{s \in S} G_{is0^-} \leq NP_i \quad \forall i \in F \quad (\text{Plane Count}) \quad (2.2)$$

$$G_{ist^-} - G_{ist^+} + \sum_{j \in \text{IN}(i,s,t)} x_{ij} - \sum_{j \in \text{OUT}(i,s,t)} x_{ij} = 0 \quad \forall i \in F, s \in S, t \in T \quad (\text{Balance}) \quad (2.3)$$

$$\sum_{i \in F} x_{ij} = 1 \quad \forall j \in J \quad (\text{Cover}) \quad (2.4)$$

$$\sum_{j \in J} R_{0jv} + \sum_{j \in J} \lambda_{jv} \left(\sum_{i \in F} \text{CAP}_{ij} x_{ij} \right) - R_{\text{Total}} \geq 0 \quad \forall v \in V \quad (\text{Revenue}) \quad (2.23)$$

$$C_{\text{Total}} - \sum_{j \in J} \sum_{i \in F} C_{ij} x_{ij} = 0 \quad (\text{Cost}) \quad (2.24)$$

$$\begin{aligned} x_{ij} &\in \{0, 1\} \quad \forall i \in F, \quad \forall j \in J \\ G_{isj} &\geq 0 \quad \forall i \in F, \quad s \in S, \quad t \in T \end{aligned} \quad (2.5)$$

For this formulation, the objective function is modified and two new constraints are added. Constraints (2.23) represent a variation of Constraints (2.21) and allows for the incorporation of the original binary decision variable, x_{ij} . Constraint (2.24) simply redefines the total cost of the fleet assignment as a constraint. O&D FAM explicitly incorporates network effects by utilizing the bidprices provided by

solving an O&D network flow model to estimate the revenue function of FAM. The network passenger flow model represents the O&D yield management or O&D revenue management process.

Conceptually, O&D FAM is very different than leg-based FAM. The revenue estimates for leg-based FAM are made on a leg-by-leg basis. Therefore, the revenue estimates and subsequent FAM formulation do not capture the network effects due to O&D yield management. Even using a nonlinear approximation to the revenue function for each leg in the network, leg-based FAM cannot accurately approximate the impact of up-line and down-line capacity constraints. On the other hand, O&D FAM incorporates the network effects directly into FAM through Constraints (2.23). Constraints (2.23) represent an upper bound on the *total system revenue* for the network and provide a link between fleet assignment and network flow as a function of the bidprices for all the legs within the network. As a result, the approximation is also a function of the O&D revenue management effects throughout the network.

For O&D FAM formulation presented above, the O&D revenue Management or O&D Yield Management (O&D YM) process is modeled using a nonlinear network flow model that maximizes overall system expected revenue subject to capacity constraints for each flight in the network. The decision variable represents the number of seats allocated to each O&D fare class itinerary. The overall expected revenue is based on the expected traffic for each O&D itinerary and is a function of the seats allocated to the itineraries competing for space over each flight leg in the network. The nonlinear O&D YM model is typically solved using a sub-gradient algorithm. Appendix A presents a complete formulation and review of the O&D revenue management problem.

To accurately implement O&D FAM, several network flow solutions corresponding to feasible fleet assignments are needed. One possible iterative approach for solving this problem is shown in Fig. 2.8. To begin the iterative algorithm, initial fleet assignments or bidprices for each leg in the network are needed. These can be obtained by: (1) assuming arbitrary initial capacities for each leg in the schedule and solving the network flow model or (2) assuming some initial bidprices and solving O&D FAM or (3) using a standard leg-FAM solution to define the initial capacities. Either option is acceptable and will not affect the final solution of the algorithm. For the work presented here, we assume initial capacities for each leg in the schedule and start the algorithm by solving the network flow model. For each iteration of the algorithm shown in Fig. 2.8, the linear FAM model is subjected to a subset of the revenue Constraints (2.23). To define the parameters needed for Constraints (2.23), the total O&D revenue ($R_{\text{Total YM}}$) is estimated using the Network Flow (O&D revenue management) model.

By incorporating O&D passenger flow aspects into the model, O&D FAM avoids many of the problems inherent to Leg-FAM. Using the results of the network flow model, O&D FAM provides a more realistic estimation of the *total system revenue* function for the network schedule. In addition, O&D FAM incorporates the impact of traffic flow in a manner that is consistent with revenue management practices. This approach uses a series of linear approximations based

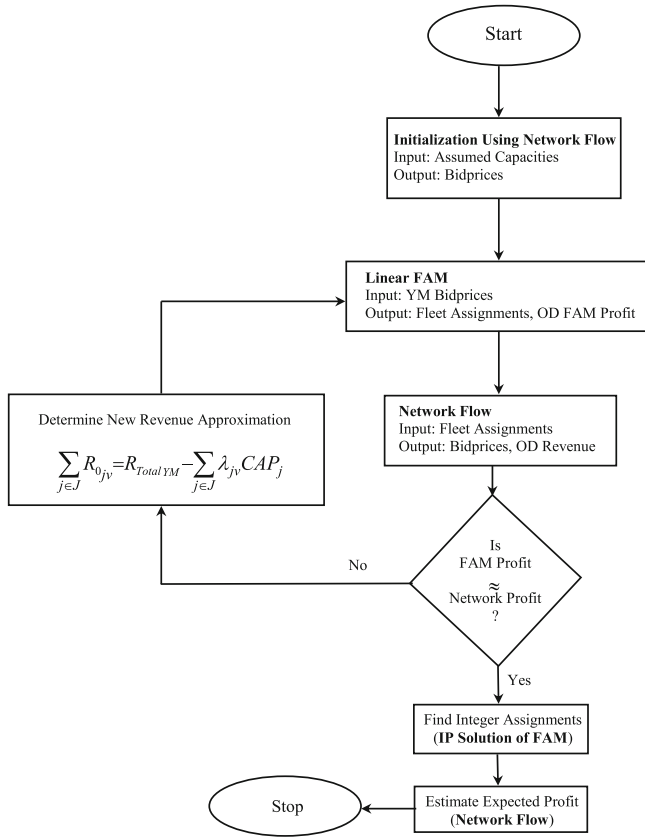


Fig. 2.8 Basic O&D FAM algorithm

on the expected network traffic to determine a linear upper bound of the total revenue function instead of point estimates of the total expected revenue for each leg. This approach prevents the propagation of errors and possible inferior fleet assignments due to inaccurate estimates of the revenues based on leg-level estimates of demand instead of the O&D traffic throughout the network. Although the decision variable is defined as a binary variable, it is computationally beneficial to relax this condition throughout the iterative portion of the solution process. This allows fractional solutions to be used in approximating the revenue function.

Although the iterative approach presented in Fig. 2.8 provides a more accurate estimation of the revenue and can result in a more efficient fleet assignment, it does require multiple solutions to the relaxed fleet assignment model. This problem is compounded by the fact that accurately modeling the revenue function may require several linear approximations of the revenue function. This can become extremely inefficient and computationally cumbersome. To avoid this problem, we can use convex combinations to generate additional bidprices and linear

approximations of the true revenue function without resolving the complete FAM formulation. Jacobs et al. (1999, 2008) provides a detailed description of how to improve the computational efficiency of this formulation using convex combinations.

Now that we have developed a conventional FAM and developed two enhancements to that model to incorporate the network effects, we turn our focus to applying these models to the fleetings of actual airline schedules. For illustration, we will consider the use of these models for general fleetings as would be used in long-term schedule planning.

To illustrate the conventional Leg-FAM and O&D FAM approaches, a ten-city, 48 leg example consisting of 534 O&D service markets and two fleets is presented and compared to the results obtained using two different revenue estimation schemes for the leg-based FAM formulation. The first leg-based formulation uses a prorating scheme to approximate the expected revenues from each leg in the schedule. The second formulation simply uses the total revenue for each O&D market class to approximate the revenue for each leg. Results indicate that O&D FAM outperforms both leg-based FAM formulations. Fleet assignments from O&D FAM show a 2.8% improvement in expected profit over the results of the leg-based FAM formulations.

To compare the conventional leg-based FAM formulation and O&D FAM, we benchmark the benefits and practicality of using an O&D approach to fleet assignment. We present a case study comparing the Leg-FAM and O&D FAM approaches using actual airline schedules with more than 4,000 daily operations. For these benchmarks, we present results of a general fleet assignment process in which all the scheduled flights must be assigned an aircraft type and a schedule reduction run in which non-profitable flights can be canceled.

2.5.3 Illustrative Example

To illustrate and compare the typical FAM and O&D FAM approaches, we use the simple ten-city example shown in Fig. 2.9. This network consists of 48 flight legs, 534 O&D market classes and two fleet types. In addition to direct flights to and from DFW, the network includes six non-stop flights. The market class prices are representative of actual coach fares for each O&D reported in SABRE. Typical market class demands are assumed. For this illustrative example, the model was formulated and solved using AMPL, the commercially available LP/IP solver CPLEX and the O&D passenger flow model presented in Appendix A.

The network flow model determines the optimal network traffic and bidprices using an approach similar to the O&D revenue management model used by many airlines. Standard deviations for determining traffic and seat allocation for each market class were calculated using a constant coefficient of variation (CV). Traffic and demand were modeled using a Gamma distribution. Table 2.2 presents a summary of the data used to formulate and solve the model.




Fig. 2.9 Network for illustrative example (Jacobs, Ratliff and Smith, 2000)

Table 2.2 Summary of data for illustrative example

Model parameter	Value
Number of stations (cities)	10
Number of flight legs	48
O&D market classes	534
Number of hub (DFW) complexes	2
Fleet types:	2
Boeing 757	
Capacity	200
Cost (\$/min)	\$79.00
Available aircraft	7
MD Super 80	
Capacity(No. of passengers)	150
Cost (\$/min)	\$66.45
Available aircraft	8
Average O&D demand	
Y Class	5
M Class	15
Q Class	25
Coefficient of variation (CV)	0.35

To compare with the O&D FAM solution, Leg-FAM must be solved using average or estimated leg-based revenues and costs. For this illustrative example,

Table 2.3 Results for prorated revenue case

(1)	Leg-FAM		O&D FAM	
	Obj. function value (2)	Expected value (3)	Obj. function value (4)	Expected value (5)
Continuous solution	\$55,202	\$66,080	\$69,886	–
Integer solution	\$55,202	\$66,080	\$69,519	\$67,954

two common and simple approaches were used to estimate the leg revenues. In the first case, the expected revenue is estimated by prorating the average fare paid by a customer proportional to the distance of each flight leg in the itinerary. In the second case, the revenues are estimated by averaging the total revenue paid by a customer. This second approach double counts the revenue and represents an extreme attempt to capture the O&D revenue. A complete description of these approaches is presented in Appendix B.

2.5.3.1 Results and Comparison with the Prorated Revenue Case

Table 2.3 compares the results from solving the ten-city example using O&D FAM with typical Leg-FAM results for the same test scenario. The results presented in Table 2.3 represent expected profit. For the results presented here, the Leg-FAM formulation used prorated leg revenues. The prorated revenues used in this case were proportional to the block times for each leg in the itinerary.

Column 2 of Table 2.3 represents the expected profit from Leg-FAM. For this example, the initial Leg-FAM solution was integer and no branch and bound was necessary. To estimate the total expected profit on an O&D level, the Leg-FAM solution is used as input to the network flow model (see Appendix A). The network flow model estimates the optimal traffic and revenue by intelligently allocating space to competing O&D fare classes. Column 3 represents the total expected O&D profit using the fleet assignments from Leg-FAM.

Column 4 presents the objective function value for the O&D FAM solution. The first entry in this column represents the continuous solution for O&D FAM. The solution to this illustrative problem required 5 iterations of the main algorithm and 2 iterations of the convex combination algorithm. The second entry represents the expected profit when the O&D FAM model is forced to an integer solution subject to the Benders approximations to the revenue function. As expected when the problem is further constrained, the total revenue decreases. This example required 19 branch and bound nodes to determine the optimal integer solution. Because this solution is constrained by the linear approximations to the total revenue, it is an over estimate of the true profit. Column 5 presents the expected profit for the integer O&D FAM solution. For this example, the O&D FAM model required eight iterations to converge to a 1% difference in the linear FAM Profit

Table 2.4 Solution comparison for flight leg JFKDFW1 total leg demand = 179.5

(1)	Leg-FAM Inputs		Network solution (O&D FAM)	Network solution (Leg-FAM)
	S80 (2)	757 (3)	S80 (4)	757 (5)
Traffic	138.9	165.4	122.6	153
Revenue per passenger	\$152.55	\$152.55	\$193.83	\$166.99
Load factor	0.93	0.83	0.82	0.77
E(revenue)	\$21,186	\$25,232	\$23,764	\$25,550
Leg cost	\$11,961	\$14,220	\$11,961	\$14,220
E(profit)	\$9,226	\$11,012	\$11,803	\$11,330

and the Network Profit. This performance varies based on the specific attributes of the problem and the convergence tolerance. The figure in column 5 represents the expected profit found by solving the network flow model for the final integer assignments from O&D FAM.

In the scenario presented here, O&D FAM shows a 2.8% improvement in total expected profit over the Leg-FAM solution. For this example, the iterative O&D FAM algorithm required approximately 20 s of real time to converge. The solution time for Leg-FAM was nearly instantaneous.

Appendix C presents the final fleet assignments and expected traffic for the O&D FAM and Leg-FAM solutions. The average overall network load factor for these results equals 66.9% for the Leg-FAM solution and 66.2% for the O&D FAM solution.

The results presented in Table 2.3 illustrate that O&D FAM can produce superior fleet assignments. However, to clearly understand the fundamental differences between the Leg-FAM and O&D FAM solutions, it is necessary to look at the solution at the leg level. Table 2.4 presents specific details a leg in which the optimal aircraft assignment differs for the two methods. In addition, this table provides details concerning the expected revenue used as input for the Leg-FAM formulation.

Table 2.4 compares the results for the JFKDFW1 flight leg in the network. For this example, Leg-FAM assigned a Boeing 757 and O&D FAM assigned a S80 to the flight leg. This flight leg serves 27 individual market classes and has a total average leg demand of 179.5 passengers. Columns 2 and 3 of Table 2.4 present the expected traffic, revenues and profit figures for use as input to Leg-FAM. The traffic values presented in columns 2 and 3 are determined using a normally distributed spill model with a coefficient of variation (CV) of 0.35. The revenue per passenger represents a weighted average based on demand of the prorated leg revenues for each market class itinerary traveling on the leg. The leg cost is the cost of flying the leg and is based on the block times for the individual flight leg. Column 4 presents the expected traffic, revenues and profits for the final O&D FAM solution. Column 5 presents the expected traffic, revenues and profit figures for the final Leg-FAM solution. For columns 4 and 5, the average revenue per

Table 2.5 Solution comparison for flight leg ORDDFW2 total leg demand = 168

(1)	Leg-FAM Inputs		Network solution (O&D FAM)	Network solution (Leg-FAM)
	S80 (2)	757 (3)	S80 (4)	757 (5)
Traffic	135.7	158.3	130.3	155.3
Revenue per passenger)	\$77.23	\$77.23	\$83.97	\$78.05
Load factor	0.90	0.79	0.87	0.78
E(revenue)	\$10,479	\$12,222	\$10,945	\$12,123
Leg cost	\$9,303	\$11,060	\$9,303	\$11,060
E(profit)	\$1,176	\$1,162	\$1,642	\$1,063

passenger represents the quotient of the total expected revenue and the expected traffic.

The results presented in Table 2.4 illustrate two consistent differences between O&D FAM and Leg-FAM solutions. First, Leg-FAM consistently *underestimates* the average revenue per passenger. Second, Leg-FAM *overestimates* the total expected traffic resulting in a load factor that is too high. These differences are due to the fact that Leg-FAM uses weighted average revenue based on total leg demand to estimate the total revenues for each leg. The spill model used to estimate traffic for Leg-FAM considers only the total demand for the leg. No consideration is given to the allocation of seats to different O&D market fare classes. In essence, using total demand to estimate the traffic assumes that all the passengers are alike and have the same chance of obtaining a seat. In reality, seats are allocated based on price, expected demand and overall revenue. The network flow model tends to reduce leg traffic and increase the average revenue per passenger by discriminating between the individual O&D market classes bidding for space on the flight leg. The total revenue received for a network schedule is a function of the individual fares and actual traffic for each O&D market class throughout the network.

The concept of leg-based revenue is misleading. For example, the traffic for the JFKDFW1 flight leg presented in Table 2.5 consists of 27 individual market classes. Only three market classes using this flight leg were local with itineraries consisting of a single flight leg. The remaining traffic traveling on this leg had itineraries consisting of two legs. These inconsistencies between the network flow solution and capacity planning lead to less profitable fleet assignments. Table 2.5 presents results for a second leg in the example network.

For Leg-FAM to yield accurate and meaningful results, the input assumptions must reflect the expected results. The results presented in Tables 2.4 and 2.5 show that the inputs to Leg-FAM differ significantly from the results with respect to revenue per passenger, load factor and total revenue.

Table 2.6 Results for total revenue case

(1)	Leg-FAM		O&D FAM	
	Obj. function value (2)	Expected value (3)	Obj. function value (4)	Expected value (5)
Continuous solution	\$543,841	\$66,080	\$69,886	–
Integer solution	\$543,841	\$66,080	\$69,519	\$67,954

2.5.3.2 Results of Total Revenue Case

To fully appreciate the benefits of O&D FAM, the results are compared to a Leg-FAM solution in which total O&D revenue is used to estimate the leg revenue. Table 2.6 compares the results from solving the ten-city example using O&D FAM with results obtained using Leg-FAM. As before, the results presented represent profits. For the results presented here, the Leg-FAM formulation used average total itinerary revenues to estimate the expected revenue for each leg in the network. The average itinerary revenues used in this case were weighted proportional to the itinerary demand. The expected revenue for each leg represents the product of the expected traffic and the average total itinerary fare paid by an individual traveling on the leg. As in the prorated case, the initial solution to Leg-FAM was integer.

In this case, the Leg-FAM solution using total revenue yields the same fleet assignment as the solution using prorated revenues. For this example network, the solution indicates that there is no benefit from using prorated revenues over total revenues in the objective function of Leg-FAM. However, the results presented in Table 2.6 do show that using total revenues to formulate Leg-FAM results in a substantial overestimate of the total expected revenue and profit. In this case, the Leg-FAM solution overestimates the *total system revenue* by 73%. This results in a Leg-FAM profit that is eight times the actual expected profit.

Since the total revenue example uses an average of the total revenue for itineraries traveling on any given leg in the network and double counts revenue for connecting passengers, it is difficult to compare individual leg assignments and draw any conclusions about the expected revenue per passenger. This type of solution only allows network-wide comparisons of the overall profit.

The illustrative examples presented above illustrate the use of both Leg-FAM and O&D FAM for fleeting airline schedules. In the ten-station, 48 flight leg example presented above, O&D FAM outperformed Leg-FAM by 2.8% in profit. This scenario consisted of 534 O&D market classes with three fare classes per O&D and two equipment types. Fleet assignments from O&D FAM show significant improvement in expected overall profit over Leg-FAM. The algorithm is computationally efficient and an integer solution for this example can be found by solving relatively few branch and bound sub-problems. Although this example is relatively small, it provides an initial benchmark for assessing the effectiveness of

the O&D FAM algorithm. This formulation of O&D FAM provides a solid platform to develop and test more sophisticated models. For this formulation of O&D FAM, all nonlinear aspects of the problem are isolated in network flow model. The O&D market effects captured by the network flow model are transferred to FAM in the form of linear approximations to the total revenue function. The fact that the nonlinearities associated with the O&D passenger flow within O&D FAM are isolated in the network flow model presents opportunities for enhancing the model to include additional aspects important to the fleet assignment and schedule planning process. These include: price optimization, passenger recapture and added crew and station fixed and manning costs in the fleet assignment process.

Next we need to apply these same models to actual airline schedules to illustrate their effectiveness, scalability and practicality. In the next section, we present a long-term planning case in which Leg-FAM and O&D FAM are used to fleet an actual airline schedule.

2.5.4 Application in Practice: General Fleet Assignment for Schedule Planning

To illustrate the benefits associated with the O&D FAM approach presented earlier, we consider the assignment of fleets to a commercial airline planning schedule and compare the expected profitability to the fleet assignments using the typical leg-FAM approach. For this illustration, we consider two types of fleet assignment scenarios. The first scenario represents a switching run in which all the flight legs in the schedule must be assigned a fleet. This type of scenario represents the typical resource allocation process used by many airlines to finalize their fleet assignments during the scheduling cycle. The second scenario represents a reduction run in which we relax the cover Constraints (2.4) to allow cancellation of unprofitable flying. This represents the process used by many airlines to prune a draft schedule during the long-range planning cycle.

The passenger revenue, traffic and costs used for this application were forecasted by American Airlines' proprietary Integrated Forecasting System (IFS). IFS forecasts market share, passenger demand, spill, traffic, revenue and cost (both variable and fixed) at either the O&D itinerary or leg level based on the host airline schedule and OA competitive schedules. In addition, IFS takes into account numerous attributes such as passenger recapture, fuel burn and cost, station and crew manning costs, minimum revenue guarantees (MRG), frequent flyer program (FFP) revenue allocations, code-share and proviso agreements and revenue accounting allocations more representative of the actual business process. For Leg-FAM, IFS allocates passenger revenue to each leg based on local and connecting passenger demand and spill, the up-line and down-line contribution of connecting passengers to the network. For O&D FAM, IFS forecasts passenger demand and revenue at the O&D itinerary level. For this application of Leg-FAM and O&D

Table 2.7 Comparison of O&D FAM versus Leg-FAM for switching scenarios

Scenario	Daily operations	Change in passenger traffic (%)	Change in revenue (%)	Change in cost (%)	Change in profit (%)
Fall	4,034	(3.65)	(1.80)	(2.40)	1.21
Spring	4,434	0.10	0.71	(1.20)	7.20

Table 2.8 Comparison of Leg-FAM and O&D FAM schedule reduction results

Fleeting scenario	Daily operations	Change in passenger traffic (%)	Change in revenue (%)	Change in cost (%)	Change in profit (%)
Base	4,930	–	–	–	–
Leg-FAM	4,569	(1.9)	(2.5)	(5.2)	11.6
O&D FAM	4,281	(13.1)	(9.5)	(14.9)	18.1

FAM, IFS was used to forecast the inputs to both models and evaluate the resulting schedules produced by each of the models

In the first case, we present the results of a switching run on two different planning schedules. Table 2.7 presents the results of the switching scenario.

The switching results presented in Table 2.7 illustrate the economic improvements associated with using an O&D FAM approach over the traditional Leg-FAM approach. For the fall schedule we see that the O&D FAM approach results in a reduction in system traffic and revenue along with an even bigger reduction in overall system costs. This resulted in an overall system profit increase of 1.21% (0.2% of revenue gain). For the spring schedule, the O&D FAM approach increased the overall system traffic and revenue while decreasing the overall system costs. This resulted in an overall profit increase of 7.20% (1.1% of revenue gain). For the switching scenarios presented in Table 2.7, O&D FAM required 10–12 iterations of the Benders algorithm to converge.

In the second case, we present the results of a reduction run which allows the cancellation of unprofitable flying. Table 2.8 presents the results of the reduction scenario.

The results presented in Table 2.8 clearly show an improvement in overall profitability for the schedule reduction scenario. Traditional leg-FAM reduced the scheduled operations by nearly 2% while O&D FAM reduced the scheduled operations by approximately 13%. In this case the overall profitability increase resulted from the model's ability to decrease costs by dropping unprofitable flying and reducing aircraft gauge significantly more than revenue reduction suffered by reducing the overall O&D services available.

In both the switching and reduction runs presented above, O&D FAM tends to reduce the overall system costs significantly more than the traditional leg-FAM approach. This is due to the fact that O&D FAM explicitly models the O&D passenger flows within the assignment process. As a result, O&D FAM properly spills low value O&D connecting traffic in favor of cost savings while leg-FAM

assigns aircraft based on a leg-centric view of profitability in which the traffic and revenue on each leg is assumed to be independent of the capacity of other flights within the schedule. This means that leg-FAM assumes the traffic and revenue for a flight leg remains constant even if the flights used by connecting passengers are down-graded or canceled. However, O&D FAM improves on the traditional leg-based methodology by incorporating the O&D passenger flow into the assignment process. As a result, passenger spill happens at the O&D level and the corresponding traffic and revenue automatically adjust as capacity assignments change or flights are canceled.

In the spring case where demand is high, O&D FAM repositions capacity such that more profitable passengers are allocated capacity while still reducing overall system costs. In these cases, O&D FAM properly balances the revenue benefits of local and connecting traffic with the potential reduction in costs. This yields a significant improvement in overall system profit.

The benchmark results presented here further illustrate the benefits associated with using an O&D approach to the fleet assignment problem. In these cases, both the switching and reduction scenarios showed a significant improvement in overall system profitability. In the switching cases, O&D FAM found solutions that significantly cut overall system costs while either marginally decreasing revenue by spilling low value connecting passengers or increasing overall system revenue by properly allocating capacity to accommodate profitable passengers.

2.5.5 Section Summary and Conclusion

This section has presented a brief overview of the scheduling and fleet assignment process. In addition, this section has provided an introduction to some of the optimization models commonly used to develop and finalize an airline schedule prior to publication and shown how these same processes and models can be used in the near-term to improve schedules to better match overall demand to allocated capacity. In the next section, we introduce some of the concepts models associated with routing a final schedule to meet crew and maintenance and engineering requirements.

2.6 Aircraft Rotations and Maintenance Routing Planning

Once the fleet assignment process has been completed, schedule planners focus on evaluating potential rotations and routing solutions for the aircraft used to fly the schedule. Aircraft rotations or turns link arriving flights to departing flights within a station while routings consider the complete path of an aircraft over the day, week or month. The aircraft rotation solution is governed by local operational requirements, operational efficiency and aircraft count. The aircraft routing process

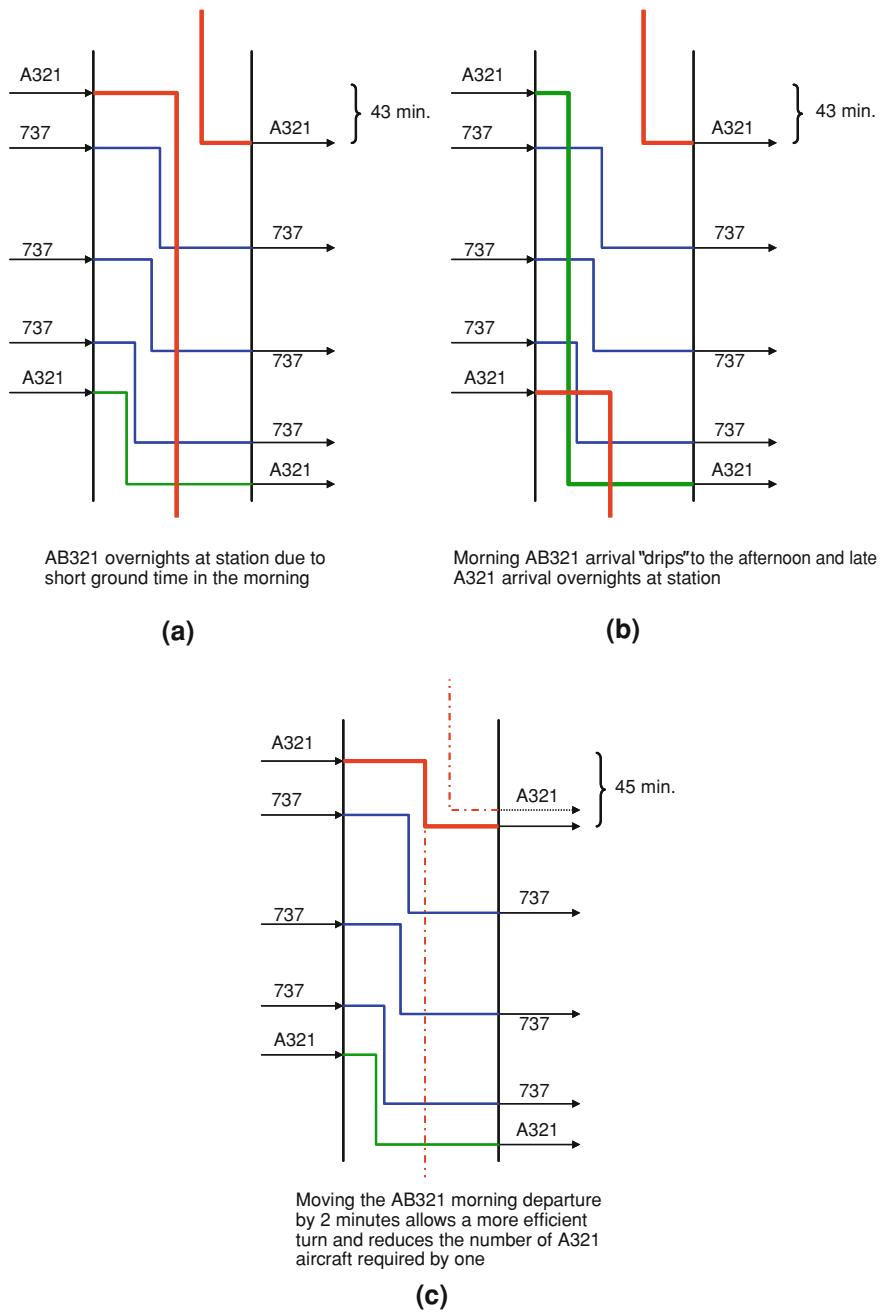


Fig. 2.10 Impact of aircraft turns. **a** AB321 overnights at station due to short ground time in the morning43. **b** Morning AB321 arrival “drips” to the afternoon and late A321 arrival overnights at station. **c** Moving the AB321 morning departure by 2 minutes allows a more efficient turn and reduces the number of A321 aircraft required by one

involves multiple levels of detail and sophistication and is usually governed by maintenance requirements. In addition, aircraft routings typically provide maintenance opportunities rather than actual maintenance scheduling for specific aircraft. The following sections provide an overview of some of the ways airlines determine rotations and evaluate routings with the schedule.

2.6.1 Aircraft Rotations

Generally, the FAM and the fleet assignment process are driven by overall system profit or cost and do not determine aircraft turns within each station in the network. As a result, the schedule planners must determine the best aircraft rotations by “connecting” or “turning” the schedule in an efficient manner. During this process, schedule planners will attempt to maximize the efficiency of the station by effectively managing ground time and removing any excessively long ground times known as “drips”. The set of rotations presented in Fig. 2.10a and b show a typical set of aircraft rotations resulting from a fleet assignment solution and applying a minimum ground time of 45 min at the station. In Fig. 2.10a, the early morning A321 arrival cannot turn to the morning A321 departure because the available ground time of 43 min violates the minimum ground time for the station. As a result, the aircraft sits at the station until the next morning when it supports the A321 morning departure. An alternative rotation pattern is shown in Fig. 2.10b. In this case, the morning A321 arrival connects to the last departure of the day. Again, one aircraft sits on the ground all day and the last arrival of the day sits overnight to support the morning departure. Both of these patterns represent inefficient rotations and increase the number of aircraft required to fly the schedule without any revenue benefit.

Moving the morning departure by 2 min improves the rotation pattern and reduces the number of aircraft required to fly the schedule by one (Fig. 2.10c). In cases where the aircraft count constraints are not tight, the FAM can produce fleet assignments that maximize overall profit but yield inefficient operations at the station level. This example only represents one of many possible scenarios uncovered by the schedule planners following the fleet assignment process.

Unfortunately many commercial carriers use simple last-in/first-out (LIFO) and first-in/first-out (FIFO) strategies to rotate aircraft. These strategies typically require a great deal of manual intervention by the schedule planners to remove inefficiencies and insure the final schedule meets all operational constraints and adheres to aircraft count limitations.

The process of determining the most efficient aircraft rotations can also be formulated as an optimization model in which the objective represents the minimization of excess or unnecessary ground time within the schedule. The decision variables reflect the aircraft connections or rotations for each station and the final timing of the flights.

Excess ground time represents the extra time an aircraft spends on the ground between flights compared to minimum ground time (MGT) needed to turn or rotate the aircraft for the next flight. When it occurs during the day, it signifies a lower utilization of the aircraft. In the early schedule planning phase, excessive ground time means extra aircraft must be used to cover the scheduled flights. Later in the scheduling process when additional operational considerations such as crew with equipment and updated gating availability limitations become a priority, the model can be used to effectively manage ground time to improve operations.

Conceptually, effectively managing aircraft rotations and ground time within the schedule requires a model that explicitly represents connections between arriving and departing aircraft and Terminator to Originator (T to O) connections that connect some arrivals to departures the following day. In addition, the model must allow limited flight retiming to facilitate improved rotations that reduce excess ground time while constrained by aircraft count.

One possible model formulation to determine efficient aircraft rotations while managing ground time within the schedule is presented below. This formulation uses the following notation:

Decision Variables

F_{jk} binary variable for flight j , retime candidate k

C_{jk_1, lk_2}^s binary variable that indicates flight j candidate k_1 is connected to flight l candidate k_2 at station s

T_{sjk} binary variable, is 1 if flight j candidate k is selected as a terminator, overnight at station s .

O_{sjk} binary variable, is 1 if flight j candidate k is selected as an originator, departing from station s

TO_{shf} integer variable, signifies overnight arc at station s for station h and assigned fleet f

Sets and Parameters

a_{sj}^{in} mapping parameter is 1 if flight j is an inbound flight for station s

a_{sl}^{out} mapping parameter is 1 if flight l is an outbound flight for station s

b_{sjk_1} mapping parameter is 1 if flight j candidate k_1 can become a terminator at station s

b_{slk_2} mapping parameter is 1 if flight l candidate k_2 can become an originator at station s

h_{sjhf} mapping parameter, is 1 if flight j is O/T at station s but is confined to hub-isolation hub h , and assigned fleet f

p_{jk_1, lk_2} penalty associated for connecting flight j candidate k_1 and flight l candidate k_2 (=square of GT)

p_f incentive associated with freeing one extra plane from fleet f

p_k penalty associated with flight j candidate k

p_{sjk} penalty associated with originator/terminators on station s and flight j

FL set of all flights considered




Fig. 2.11 Ground time management model arc formulation

S set of all stations considered

H set of all hubs considered for hub isolation (should include free-hub + 1 hub iso, like ORD)

$K(j)$ set of retime candidate for flight j

$E(j, k_1)$ set of outbound flights which is allowed to connect to flight j candidate k_1

$E(l, k_2)$ set of inbound flights which is allowed to connect to flight l candidate k_2

F set of fleet

Using these definitions, the problem of determining efficient aircraft rotations while managing excess ground time and aircraft count is formulated as:

$$\begin{aligned}
 \text{Minimize } & \sum_{l \in E(j, k_1)} p_{jk_1, lk} C_{jk_1, lk}^s - \sum_{f \in F} p_f s_f + \sum_{j \in F} \sum_{k \in K(j)} p_k F_{jk} + \sum p_{sjk} O_{sjk_1} \\
 & + \sum p_{sjk} T_{sjk_1}
 \end{aligned} \tag{2.25}$$




Fig. 2.12 Impact of various re-timing windows on aircraft count

Subject to:

$$\sum_{k \in K(j)} F_{jk} = 1, \quad \forall j \in FL \quad (2.26)$$

$$a_{sj}^{in} F_{jk_1} - \sum_{l \in E(j, k_1)} C_{jk_1, lk_2}^s - b_{sjk_1} O_{sjk_1} = 0, \quad \forall k_1 \in K(j), \quad \forall j \in FL \quad (2.27)$$

$$a_{sl}^{out} F_{lk_2} - \sum_{j \in E(l, k_2)} C_{jk_1, lk_2}^s - b_{slk_2} T_{slk_2} = 0, \quad \forall k_2 \in K(l), \quad \forall l \in FL \quad (2.28)$$

$$TO_{shf} - \sum_j \sum_{k_1 \in K(j)} h_{sjhf} O_{sjk_1} = 0, \quad \forall s \in S, \quad \forall h \in H, \quad \forall f \in F \quad (2.29)$$

$$TO_{shf} - \sum_j \sum_{k \in K(l)} h_{slhf} T_{slk_2} = 0, \quad \forall s \in S, \quad \forall h \in H, \quad \forall f \in F \quad (2.30)$$

$$\sum_{h \in H} \sum_{s \in S} TO_{shf} + s_f = PL_f, \quad \forall f \in F \quad (2.31)$$

Solutions to this formulation determine efficient aircraft rotations that minimize the excess ground time and releases as many aircraft as possible from the schedule. The objective function (Eq. 2.25) minimizes the penalty associated with each aircraft rotation. For this formulation, the penalty for connecting an arrival to a viable departure increases as a function of the elapsed ground time between the arrival and departure. Second term in the objective function provides an incentive for the model to release or un-assign aircraft from the schedule. The remaining

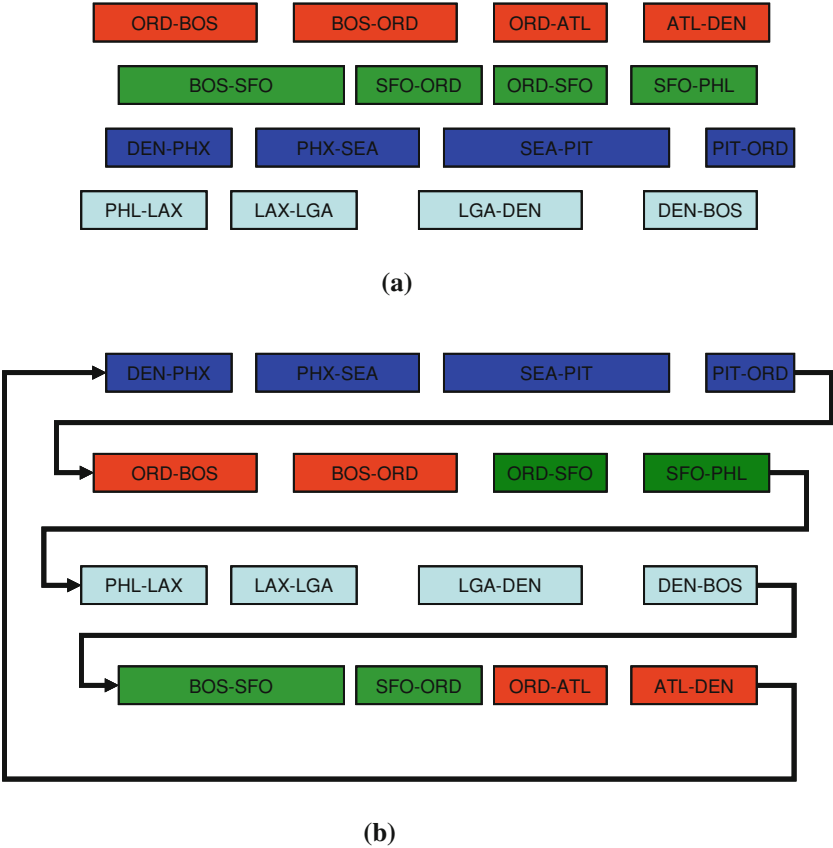


Fig. 2.13 Creating an Euler tour for aircraft routings

penalties help to limit the number of re-timed flights, terminators and originators within the finished schedule.

The model constraints, Constraints 2.2 through 2.6, describe various details of flow balance within the schedule. The first constraint, Constraints 2.2, insures that only one retime candidate be selected for each flight within the schedule. Constraints 2.3 represent the flow balance constraint for inbound flights and potential connections to departing flights. Similarly, Constraints 2.28 represents the balance constraint for the outbound flights and each potential connection to a viable inbound flight. In addition to connecting to a departing flight on the same day, a flight can also stay overnight at the station. In this case, the inbound flight connects to a terminator arc. Similarly, an originating flight will connect to an outbound flight the following day. The terminators and originators are connected via the integer T–O arc. Figure 2.11 presents an illustration of the connections for a single station assuming a 1 day schedule. Constraints 2.29 and 2.30 link the terminator

arcs with the T–O overnight arc and the T–O overnight arc to the originator arcs, respectively. Finally, Constraint 2.31 limits the number of aircraft used in the schedule to be less than or equal to the aircraft available. Because this formulation only considers aircraft connection and not fleet assignments, the aircraft count can be modeled by simply adding up the number of T–O aircraft. The slack variable in the aircraft count constraint allows the schedule planner to incentivize reducing the number of aircraft required to fly the schedule.

Figure 2.12 presents the results of running the aircraft rotation model for various retiming windows for a schedule of approximately 225 aircraft. As the results show, for small retiming windows, the model solution does not reduce the number of aircraft required to fly the schedule a great deal. However, allowing the model to use a maximum retiming window of ± 20 min can reduce the number of required aircraft by nearly 4.5% or about ten aircraft.

2.6.2 Aircraft Maintenance Opportunity Routing

The initial aircraft maintenance routing process focuses on evaluating routing feasibility of the fleet schedule to meet various maintenance and operational requirements. In reality, the schedule planners simply try to insure that aircraft have ample opportunity to receive maintenance within specified time intervals. The routing of individual aircraft within the schedule happens much closer to the day of departure and changes frequently due to operational events such as weather delays throughout the network.

FAA regulations require periodic maintenance checks for commercial aircraft based on metrics such as flown hours, the number of take-offs and landings and elapsed time in service since the last maintenance check. These safety and maintenance requirements are very strict and aircraft not meeting the minimum requirements will be grounded.

In addition to meeting regulatory requirements, many carriers try to uniformly add wear and tear to the aircraft. For example, over a period of one month it is better to add approximately 300 flying hours to all aircraft of a given fleet rather than add 400 hours to half of the fleet and only 200 hours to the remaining fleet. To accomplish this, many carriers attempt to develop routings that represent an Euler Tour. For the maintenance routing problem, the nodes represent stations or airports and the arcs represent lines of flying over the day. An Euler Tour in a directed graph represents a closed tour such that each arc or line of flying is traversed exactly once even though nodes or airports may be traversed multiple times. Once an Euler Tour exists, all aircraft will repeatedly experience the same sequence of dial routes though in any given day, each aircraft is assigned a different daily route.

Figure 2.13 presents a simple example of potential daily aircraft routings for a fleet of four aircraft. These lines of flying can be determined using a number of techniques such as LIFO (Last In-First Out) or FIFO (First In-First Out) routings.

For this example, assume Chicago (ORD) has a qualified maintenance base for the airline and that routine maintenance checks (an A-Check) must be performed once every 4 days. Figure 2.13a illustrates the original daily routings for the four aircraft.

If we link the original lines of flying by matching up the final airport in any line of flying to a common originating airport, we end up with two separate tours. The first tour links the first line of flying to the third and the second tour links the second line of flying to the fourth. This results in one tour that provides overnight access at ORD and one tour that does not because the aircraft overnights at either BOS or PHL which are not maintenance bases. These types of tours are often referred to as “locked rotations” due to the fact that the aircraft flying these rotations have no opportunity for overnight maintenance. Left alone, the two aircraft flying these locked rotations would eventually violate the mandated maintenance requirements and be grounded. In such cases, the carrier would be required to ferry the aircraft to a maintenance base for the required maintenance before resuming commercial service.

However, for our four-aircraft example, we can remove the locked rotation by making a swap when two aircraft are both on the ground at the same station. To build an Euler Tour given the lines of flying shown in Fig. 2.13a, we must break two or more of the original lines of flying. For our example, this can be done by breaking and swapping the first and second lines of flying of when both aircraft are on the ground in ORD. Making this swap allows us to link the daily routings as shown in Fig. 2.13b. As a result, the locked rotation has been removed and each of the four aircraft will spend the night at the maintenance base in ORD once every 4 days. In addition, each of the aircraft will follow the same path over the 4 day maintenance interval. By routing the aircraft this way (Fig. 2.13b), the four aircraft will all accrue roughly the same number of flying hours, take-offs and landings over time. This will help the airline control the wear and tear for the fleet over time, better plan parts inventories and allow for better long-term fleet planning.

Although the two-step algorithm outlined above only considers feasibility via maintenance opportunities and ignores cost and capacity issues, it represents a common and popular approach used by many airlines and freight carriers. For the most part, it provides the schedule planners with a degree of confidence that enough maintenance opportunities exist within the schedule to meet any regulatory needs. However, additional constraints can complicate matters. For one, many airlines have consolidated the number of maintenance bases to reduce costs. In addition, some carriers perform a variety of checks (such as A, B and heavy C checks) at a single maintenance base. As a result, many carriers currently have capacity limitations associated the number and combination of checks that can be performed on any overnight shift.

The literature is replete with heuristics and formal optimization approaches to determine feasible routing schemes for a carrier. Qi et al. (2004) present a nice overview of many of the approaches used to develop feasible aircraft maintenance routings during the schedule planning phase. Clarke et al. (1997) model the aircraft routing problem as a network flow model with side constraints. This formulation

contains both flight arcs and “waiting” arcs while the nodes represent flight departures and arrivals. In this formulation, maintenance opportunities are included within the waiting arcs with enough idle time to cover the maintenance check at a specified maintenance base.

In addition, Goplan and Talluri (1998) present a polynomial time approach for solving the aircraft routing problem when maintenance must be performed at least every 3 days. In this case, they develop an iterative two-phase algorithm that builds one-day routings in the first phase and determines a tour satisfying the maintenance requirements in the second phase. If no cycle exists, a new iteration begins by modifying the one-day routes from the first phase using a heuristic that includes changing aircraft type for specific flights. Talluri (1998) shows that expanding this work to consider the four-day or more maintenance routing problem results in an NP-complete problem in the second phase. Barnhart et al. (1998) explore the use of string models for aircraft assignment and routing. In this work, they present a daily version of the string model that contains additional connectivity constraints to ensure that the strings can be concatenated to form a feasible maintenance tour or cycle. Lacasse-Guay et al. (2010) compares and contrasts several of the available aircraft routing models available in the literature. More recently, Zhe et al. (2011) present a new compact network representation of the maintenance routing problem and develop a new mixed-integer linear programming formulation to solve the problem. The proposed method can be integrated with other aspects of the airline planning process such as fleet assignment and crew pairing. In addition, Gronkvist (2005) has proposed a column generation—constraint programming hybrid approach for scheduling/routing aircraft while making tail assignments.

To develop planned maintenance routings, many carriers have implemented commercially available software applications such as Sabre’s ARM application to help the schedule planner make sure that the fleet schedule provides adequate maintenance opportunities every 3–4 days. Most of these applications use the two-phase approach outlined above and use a heuristic search and swap algorithms to find break daily lines of flying to develop feasible aircraft maintenance tours.

The reader should keep in mind that aircraft maintenance routing during the schedule planning process focuses on routing feasibility. No tail assignments are done at this time and often, the actual routings are planned much closer to the day of departure. In addition, many of the final tail assignments and routings are changed during the day of departure. In fact, at large carriers only about 80–85% of the expected terminators at 1,100 in the morning actually end up terminating at the planned stations that evening. Often, these changes are due to schedule disruptions due to weather events, crew misconnects or short-term maintenance issues that delay an aircraft and force a substitution or move-up (when an available aircraft is moved up to cover a delayed aircraft) in the schedule.

2.7 More Recent Developments in Schedule Planning and Development

In addition to the typical tasks associated with schedule planning and development, there have been a number of new developments that improve the schedule efficiency and profitability. These include the integration of crew and scheduling processes to develop more efficient flying schedules, the implementation of demand driven dispatch where fleet assignment changes are made between crew compatible aircraft close to the day of departure to better match passenger demand and the addition of probabilistic evaluation techniques to better evaluate the likely performance of specific markets.

2.7.1 *Integration of Crew and Scheduling Processes*

Originally, scheduling, crew planning, aircraft routing and revenue management processes were all optimized within separate functional silos. This was predominately due to the fact that each of these processes was done by hand. Later, when computers were available and mathematical models were developed to optimize aspects of the business, computational power was limited and airlines could not solve the more integrated and complex problems. However, many researchers have recently begun exploring integration ideas throughout the airline business. In addition to the integration of fleet and revenue management presented earlier, Sandhu and Klabjan (2007) propose the integration of schedule fleet, aircraft routing and crew pairing. Their approach combines these three aspects of the airline planning process to a single solution framework to develop a simultaneous solution to all three planning problems. In this work they present two different decomposition strategies to solve the integration problem: the first uses a Lagrangian relaxation and delayed column generation while the second uses a Benders decomposition strategy. Gao et al. (2009) investigate crew and fleet assignment integration using station purity. In their work, Gao et al. focus on three challenges that included the influence of fleet assignment on crew scheduling, addressing crew scheduling in a tractable way within the integrated model and to produce a robust schedule. Their work proposes a new approach that integrates crew connections within the FAM and imposes station purity by limiting the number of fleet types and crew bases allowed at each station.

More recently, Ruther (2010) has proposed a multi-commodity flow model for integrating aircraft routing, crew scheduling and tail assignment. In this work, Ruther assumes crews are initially only told when they will work. This allows the generation of a schedule much closer to the planning horizon. However, for many U.S.-based carriers, this assumption would violate existing collective bargaining agreements with crew rank and file. Ruther presents solutions for very small

sample cases but acknowledges that solving larger more realistic problems would be very difficult and would require other techniques to decompose the problem.

Recently, the integration of crew scheduling aspects has yielded significant cost savings associated with reduced crew connections where crews move from one aircraft to another following their arrival at a station and better feedback between the schedule development process and crew scheduling process. Integrating the scheduling and crew planning processes provides opportunities for having the crew and aircraft stay together as much as possible during the day. This saves connection time for the crew members during aircraft rotations and increases their efficiency for the airline. Some U.S.-based carriers have started the integration of crew and scheduling processes by refining the fleet assignments between sub-fleets to match crew pairings.

One of the easiest ways to accomplish this is to incentivize the FAM to consider potential crew trips within the fleet planning process. This can be done using forced turns or soft-forced turns in which a desired aircraft turn corresponds to a crew turn within a pairing or trip. Soft-forced turns represent a relaxation of the typical forced turn and allow a crew connection to be made at a cost to determine a feasible scheduling solution. In this case, the penalties or costs associated with creating a crew connection are included in the objective function. In reality, this represents an iterative process between the fleet planning model and the crew pairing optimizer.

American Airlines applied this strategy in 2006 and was able to reduce the number of flight deck crew connections for the Super 80 fleet from about 330 to 35 per day. In this case, the collective bargaining agreement gave pilots an extra 10 min to close out the aircraft paper work when connecting. By minimizing the number of crew connects, the airline was able to save approximately 100 hours of pilot and first officer labor per day of operation for the Super 80 fleet. Other carriers like Southwest Airlines have also had success with more closely integrating their crew routings with the aircraft to keep the crew and equipment together throughout the day.

2.7.2 Demand Driven Dispatch

Another area of great interest today focuses on the re-allocation of capacity close to the day of departure. Often referred to as Demand Driven Dispatch or D³ and first proposed by Berge and Hopperstad (1993) this process attempts to increase overall profitability by making strategic aircraft swaps between crew compatible aircraft near the day of departure driven by updated O&D passenger demand forecasts. Three primary factors driving the increased profitability include (1) daily forecast variability, (2) forecast error and (3) inconsistencies between the schedule planning and revenue management methodologies used to produce the schedule and manage the seat inventory.




Fig. 2.14 Incremental benefit of D³ as a function of time before departure

Table 2.9 Initial D³ benchmark results

Measure	Input schedule	D ³ solution
Incremental profit gain (% of total revenue)		0.64
Switched flights		114
Segments flown		
RJ3	230	198
RJ4	336	368
Utilization		
RJ3	10:31	9:37
RJ4	10:02	10:14

Table 2.10 Trade-off between revenue gain and swap limit

Swap limit	Daily profit increase (% of revenue)	Cumulative percent of total (%)
25	0.25	39
50	0.35	56
75	0.50	78
100	0.60	94
114	0.64	100




Fig. 2.15 Incremental passenger bookings after D³ equipment swaps

Most airlines develop schedules based on typical day or typical day-of-week forecast data. However, in reality, passenger demand can vary substantially by day or day-of-week throughout the schedule’s effective life span. D³ capitalizes on opportunities created by the systemic variation of O&D passenger demand flowing through the network by identifying strategic and feasible capacity realignments that increase total system profit. D³ exploits these opportunities by allocating additional seat inventory in markets with increased demand and decreasing seat availability in markets where expected demand has not materialized.

D³ also improves schedule profitability by using improved forecast data nearer the day of departure. Most airlines develop schedules far in advance of their operation based on average demand forecasts for the schedule period. Time lines for developing schedules vary by airline but typically range from 6 to 18 months prior to departure. The errors associated with the demand forecasts used by the revenue management systems decrease closer to the day of departure as bookings materialize.

Lastly, many airlines fleet the schedule using leg-based methods while managing the seat inventory using O&D-based strategies. This leads to an inconsistent matching of supply (aircraft capacity) and passenger demand. Using the O&D FAM methodology presented earlier, D³ can help compensate for the inconsistencies inherent in the schedule development and revenue management processes by better matching O&D passenger demand and capacity.

Table 2.11 Summary of traffic changes for June 22nd D³ swaps

Flight	Original fleet assignment	D ³ fleet assignment	Expected traffic (passengers)	Actual traffic (passengers)
1	EM3	EM4	21	21
2	EM3	EM4	48	44
3	EM3	EM4	27	32
4	EM3	EM4	47	33
5	EM3	EM4	43	38
6	EM3	EM4	36	25
7	EM3	EM4	23	27
8	EM3	EM4	48	50
9	EM3	EM4	37	34
10	EM3	EM4	46	42
11	EM3	EM4	45	50
12	EM3	EM\$	48	43
13	EM3	EM4	44	26
14	EM3	EM4	48	33
15	EM3	EM4	38	44
16	EM3	EM4	21	6
17	EM4	EM3	30	30
18	EM4	EM3	32	37
19	EM4	EM3	29	Canceled
20	EM4	EM3	36	Canceled
21	EM4	EM3	26	31
22	EM4	EM3	34	37
23	EM4	EM3	36	30
24	EM4	EM3	33	26

Table 2.12 Yearly summary of D³ incremental traffic gains

	Incremental passenger accommodated	Number of days D ³ implemented during the month	Average number of incremental passengers per day
January	56	13	4.31
February	61	22	2.77
March	276	28	9.86
April	124	29	4.28
May	239	31	7.71
June	263	28	9.39
July	189	22	8.59
August	315	31	10.16
September	138	20	6.9
October	427	31	13.77
November	327	30	10.90
December	394	20	19.70
<i>Yearly Total</i>	2,809	305	9.21

Each of these factors contributes to the potential profit opportunities of implementing a D^3 strategy within a traditional network carrier. To illustrate the benefits of D^3 scenario, we performed a series of benchmarks and pilot studies for American Eagle Airlines (a subsidiary of American Airlines). American Eagle operates approximately 2,200 daily flights within the United States, Mexico, Canada and the Caribbean Islands.

To illustrate the benefits of D^3 , Jacobs et al. (2001, 2008) applied D^3 to American Eagles' Embraer 135 and 145 aircraft during the spring and summer of 2001. The pilot study included 5 days of departures at Chicago's O'Hare International Airport. The D^3 swaps were identified and implemented using Sabre's commercial O&D FAM (Jacobs et al. 2000) implementation 14 days before departure. This timeframe was selected to maximize the potential profit benefits without impacting maintenance planning, operations and scheduling. Figure 2.14 presents the incremental benefit of D^3 as a function of the time before departure. Fourteen days before departure presented a convenient time to make D^3 swaps due to 14-day advanced purchase requirements. In addition, inside 7–10 days before departure, equipment swaps can disrupt maintenance planning.

Figure 2.14 also illustrates the individual components (daily variation in demand, improved forecasting and an O&D FAM) that make up the incremental benefit of D^3 swaps as a function of time.

Table 2.9 presents the overall results of the benchmark study using a daily schedule. The initial benchmark study did not include any restrictions on the number of equipment swaps made by the D^3 process. As a result, this benchmark result represents an upper bound on the profit benefits associated with this specific data. Overall, the unrestricted profit benefits associated with this demand and revenue data amount to a 0.64% margin profit gain (as a percent of total system revenue) and required 114 equipment swaps. To evaluate the sensitivity of D^3 to equipment swap limits, we ran a series of benchmark scenarios with various swap limits in O&D FAM. Table 2.10 presents the trade-off between the swap limits and the expected daily benefit.

To confirm the benchmark results, we performed a pilot study using 5 days of departures at Chicago's O'Hare International Airport (ORD). The 5 days included high, medium and low demand days. The pilot study allowed up to 25 equipment swaps and was driven by O&D revenue management data (itinerary fares, passenger demand and existing bookings). Table 2.9 presents the average daily results for the pilot study. Overall, the pilot study accommodated 33 additional passengers per day with an incremental profit of 0.27% with respect to total revenue. Figure 2.15 presents the cumulative number of additional passengers booked with respect to days before departure. Table 2.11 summarizes the change in traffic for D^3 swapped flights for June 22nd departures.

Following the benchmark and pilot studies, American Eagle implemented the D^3 process using O&D FAM for operations in its Chicago (ORD) and Dallas–Fort Worth (DFW) hubs. Table 2.12 summarizes the results for a full year of D^3 operations in Chicago (ORD) using Embraer 140 and Embraer 145 aircraft with a six-seat difference and a swap limit of 25. A swap limit of 25 generally constrains

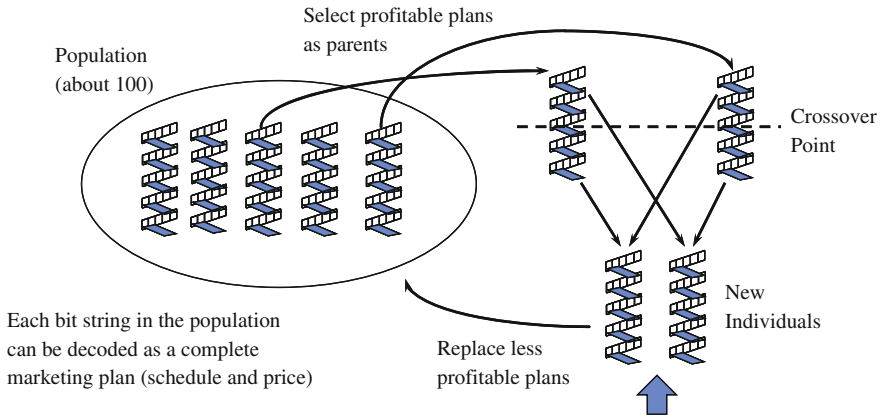


Fig. 2.16 Illustration of the genetic algorithm proposed by Walker (2001)

the number of upgraded flights to 12 per day. In addition, most of the Embraer flying out of ORD and DFW represents “out and back” round trips which mean only six upgraded flights drive the D^3 swaps.

The first column lists the month of operation, the second column represents the number of incremental passengers above the initial aircraft capacity gained by the D^3 process, the third column indicated the total number of days D^3 was used for each month and the fourth column represents the average number of passengers gained by the D^3 process per day for each month.

For this application of D^3 , results show that adjusting fleet decisions closer to the day of departure clearly identifies capacity changes that increase the O&D passenger revenue and profit in a manner consistent with revenue management controls. The D^3 swaps provided substantial incremental revenue without hampering operations at either hub.

2.7.3 Operational and Market Performance Evaluation

Another recent focus area involves the probabilistic evaluation of schedule performance. Several researchers and technology companies have developed ways to evaluate the profitability and operational performance of a proposed schedule. Rosenberger et al. (2002) and Lee et al. (2003) have developed and tested an operational performance simulation routine called SimAir to measure airline performance. This approach uses a modular set of simulation routines that reflect various aspects of the airline operations. This model also includes the ability to simulate the impact of disruptions to the airline due to inclement weather events or other unplanned delays. Green (2002) applied SimAir to evaluate the impact of various disruption scenarios on the overall operational performance of an




Fig. 2.17 GA performance with respect to network profit

American Airlines schedule. The output of SimAir consists of estimates of many of the operational metrics used by most airlines and the U.S. Department of Transportation (DOT). These include D0, A+14, completion factor and block performance.

Sabre and a number of airlines have combined efforts to develop a revenue simulation routine that can be used to estimate the impact of various revenue management strategies on overall network revenue (Ratliff 2008). This simulation routine provides participating airlines with a flexible mechanism to plug in and test various revenue management and pricing ideas and measure the impact to overall system revenue. The simulation modules use Bayesian updating scheme to develop booking scenarios that serve as inputs to a Monte Carlo simulation engine.

Wong and Jacobs (2009) and Jacobs (2005) present a risk-based methodology for quickly evaluating likely market performance of individual markets using estimates of revenue and cost factors. This methodology was encouraged by the fact that many carriers use simple estimates of revenue per available seat mile and cost per available seat mile to estimate if a market will yield profitable results without accounting for the potential variability associated with factors like market changes, operational disruptions and demand fluctuations. This approach uses the first and second moments of the revenue and cost factors to estimate the probability that the overall revenue will exceed the overall cost within the market. This methodology provides an attractive and easy mechanism for incorporating chance constraints into optimization models such as the FAM to help guide scheduling decisions based on probabilistic considerations. Parker (2008) has enhanced this approach to associate a confidence level with market share estimations and passenger behavior.

2.7.4 Automated Schedule Creation Methodologies

Ideally, schedulers would like to have a method for automatically developing a “clean-sheet” schedule that includes pricing points and accurately incorporates the revenue management, operational and cost factors into the process. Walker (2001) presents an interesting Genetic Algorithm (GA) approach for developing a “clean-sheet” schedule or incrementally improving an existing schedule. In this work, the GA represents a complete market plan (schedule and price) as a bit string. These bit strings or market plans are then combined and mutated to produce a series of new market plans that replace existing plans that represent inferior solutions for the market. This process is repeated for a set number of iterations to develop an improved overall schedule that includes both local and flow markets. Figure 2.16 illustrates this basic concept.

For the GA approach, Walker uses a simple passenger choice model to determine the attractiveness of a service within the market plan. Using this, the expected traffic is determined using the same revenue management model used in the development of the O&D FAM formulation presented earlier (See Appendix for a detailed description of the model). Costs were estimated as a function of the flights based on equipment, distance and the appropriate unit costs. Figure 2.17 presents a graphical representation of the GA performance with respect to network profitability for a small illustrative schedule design problem.

Appendix A: Network Flow and the O&D Yield Management Formulation

The O&D yield management (O&D YM) problem is a nonlinear network flow problem. The formal decision variable represents the number of seats allocated to a specific O&D market class. The objective is to maximize the overall revenue of the nonlinear objective function. The model is subject to one structural constraint that limits the total number of seats allocated to all market classes across any flight leg in the schedule to be less than the capacity of the aircraft assigned to the leg. Formally, the O&D YM model is defined by:

$$\max R_{\text{Total}_{YM}} = \sum_{s \in S} R_s \left[\int_0^{\hat{a}_s} y_s f(y_s) dy_s + \hat{a}_s \int_{\hat{a}_s}^{\infty} f(y_s) dy_s \right] \quad (\text{A.1})$$

subject to:

$$\sum_{s \in S(j)} \hat{a}_s \leq \text{CAP}_j \quad \forall j \in J \quad (\text{A.2})$$

$$\hat{a}_s \geq 0 \quad \forall s \in S \quad (\text{A.3})$$

where \hat{a}_s is the decision variable which represents the number of seats allocated to service s contained in the overall set of services S . $S(j)$ represents the set of services that involve flight j and is a subset of the total services S . A service represents a path specific O&D market fare class and can involve one or more flight legs. For example, all Y class passengers flying from Seattle to Boston through DFW on given flights are considered a single service. CAP_j represents the capacity of the fleet type assigned to leg j . y_s represents the stochastic demand for service s and $f(y_s)$ is the probability density function of the service demand. R_s equals the revenue per passenger resulting from the sale of a seat for service s . The probability density function for this formulation is usually modeled using a Gamma function. The O&D YM formulation is solved using Lagrangian relaxation (Reeves 1993) and the sub-gradient algorithm.

The solution of this model results in a set of seat allocations for each O&D service on each flight leg in the schedule. More importantly, the dual variables associated with each flight leg represent the marginal value of having one additional seat for the flight leg. This value is commonly referred to as the bidprice or hurdle rate and represents the slope of the revenue function at that leg capacity. In addition, the bidprice can be used as a metric to determine if it is profitable to sell additional tickets for a specific service.

For example, if the bidprice for traveling from Seattle to DFW is \$123 and the cost of a discount ticket for that service is \$90.00, it would not be profitable to sell the seat to a discount passenger. On the other hand, if the full fare coach ticket costs \$130, it would be profitable to sell additional full fare coach tickets.

Appendix B: Estimating Leg Revenue in FAM

One of the most common ways to formulate a FAM is using average or estimated leg-based revenues and costs. The leg-based revenues represent point estimates of the expected revenues over each leg in the schedule. Typically, there are two ways of estimating the expected revenues for each leg.

The first approach defines the expected revenues for each leg as the product of the expected traffic and the average fare paid by an individual traveling on the leg. There are two major problems with this approach: double counting of the through passenger traffic and the averaging of the revenue per passenger. This method of estimating revenues results in total expected revenues much higher than the revenues actually realized by the schedule. This is due to the fact that this approach double counts revenue by using O&D fares for connecting passengers to estimate the expected revenues for each leg in the itinerary. In addition, this approach assumes that all the demand has a per passenger revenue equal to the mean revenue for the leg. Taken together, these shortcomings can add a significant amount of error to the expected revenues for each leg in the schedule. Further

more, this approach tends to bias leg-FAM solutions to favor the assignment and use of larger aircraft. As a result, the fleeted schedule tends to under perform with respect to expected profit.

The second approach improves upon the first by prorating the revenue estimates of through traffic with respect to factors such as the length of the individual legs within the itinerary. This approach provides a more accurate and reasonable estimate of the revenue for each leg. However, this approach still assumes that the total demand on each leg has a leg revenue per passenger equal to the mean revenue for the leg.

The first approach significantly over emphasizes the value of connecting traffic on each leg within the network. This can result in the false perception that a large capacity aircraft is needed for a moderately loaded leg *fleeting differences* in the network. The second approach over emphasizes the value of the local traffic on the network legs. This approach also may result in the false perception about the size of aircraft that needed to meet the traffic needs of the leg. Reality is somewhere in between these approaches and depends on the overall network load. In cases when the average leg revenue is consistent with the O&D revenue expected on the schedule legs, Leg-FAM provides a good O&D solution. However, this is rarely the case and the errors in the revenue estimates can lead to inferior fleet assignments.

Appendix C: Final Assignments for Leg-FAM and O&D FAM (Prorated Revenue Case) (Shaded rows represent differences between Leg-FAM and O&D FAM)

	Leg-based	FAM solution	O&D FAM	Solution
Flight Leg	Equipment	Expected traffic	Equipment	Expected traffic
ABQDFW1	757	165	757	165
ABQDFW2	757	170	757	171
AUSDFW1	757	163	757	163
AUSDFW2	757	166	757	167
BOSDCA1	S80	38	S80	38
BOSDFW1	S80	124	S80	123

(continued)

(continued)

	Leg-based	FAM solution	O&D FAM	Solution
BOSDFW2	757	153	757	147
BOSORD1	757	34	757	34
BOSSFO1	S80	34	S80	34
DCABOS1	S80	38	S80	38
DCADFW1	757	158	757	157
DCADFW2	757	161	757	159
DCAJFK1	S80	34	757	34
DFWABQ1	757	170	757	171
DFWABQ2	757	148	757	149
DFWAUS1	757	166	757	170
DFWAUS2	757	145	757	144
DFWBOS1	757	162	757	158
DFWBOS2	S80	97	S80	97
DFWDCA1	757	169	757	168
DFWDCA2	757	119	757	117
DFWJFK1	757	160	757	160
DFWJFK2	757	137	S80	112
DFWORD1	757	167	S80	131
DFWORD2	S80	124	S80	124
DFWSAT1	757	169	757	170
DFWSAT2	757	149	757	149
DFWSEA1	S80	136	S80	136
DFWSEA2	757	153	757	153
DFWSFO1	S80	131	S80	131
DFWSFO2	757	145	757	144
JFKDCA1	S80	38	757	38
JFKDFW1	757	153	S80	122
JFKDFW2	757	158	757	158
JFKSFO1	S80	34	S80	34
ORDBOS1	757	63	757	63
ORDDFW1	757	131	757	131
ORDDFW2	757	155	S80	130
ORDSEA1	S80	39	S80	39
SATDFW1	757	170	757	170
SATDFW2	757	153	757	154
SEADFW1	757	183	757	183
SEADFW2	S80	24	S80	24
SEAORD1	S80	63	757	63
SFOBOS1	S80	34	S80	34
SFODFW1	757	178	757	176
SFODFW2	S80	53	S80	53
SFOJFK1	S80	39	S80	39

References

- Abara J (1989) Applying integer linear programming to the fleet assignment problem. *Interfaces* 19(4):20–28
- Algers S, Beser M (2001) Modelling choice of flight and booking class—a study using stated preference and revealed preference data. *Int J Serv Technol Manag* 2:28–45
- Barnhart C, Boland NL, Clarke LW, Johnson EL, Nemhauser GL, Shenoi R (1998) Flight string models for aircraft fleet assignment and routing. *Transp Sci* 32(3):208–220
- Barnhart C, Kniker T, Lohatepanont M (2002) Itinerary-based fleet assignment. *Transp Sci* 36(2):199–217
- Berge ME, Hopperstad CA (1993) Demand driven dispatch: a method for dynamic aircraft capacity assignment, models and algorithms. *Oper Res* 41(1):153–168
- Bradley SP, Hax AC, Magnanti TL (1977) *Applied mathematical programming*. Addison-Wesley, Menlo Park
- Carrier E (2008) Modeling the choice of an airline itinerary and fare product using booking and seat availability data. Ph.D. Dissertation, Department of civil and environmental engineering, Massachusetts Institute of Technology, Cambridge, MA
- Civil Aeronautics Board (1970) Effect on total market traffic of changes in quality of service (QSI). Docket 21136 (box 1196). Exhibit BOR-R-300. In Exhibit Series *Rebuttal Exhibits of the Bureau of Operating Rights* by J.F. Adley and C.J. Caridi
- Clarke L, Johnson E, Nemhauser G, Zhu Z (1997) The aircraft rotation problem. *Ann Oper Res* 69:33–46
- Coldren, GM (2005) Modeling the competitive dynamic among air-travel itineraries with generalized extreme value models. Ph.D. Dissertation, Department of Civil and Environmental Engineering, Northwestern University, Evanston
- Coldren GM, Koppelman FS (2005a) Modeling the competition among air-travel itinerary shares: GEV model development. *Transp Res Part A* 39(4):345–365
- Coldren GM, Koppelman FS (2005b) Modeling the proximate covariance property of air travel itineraries along the time-of-day dimension. *Transp Res Rec* 1915:112–123
- Coldren GM, Koppelman FS, Kasturirangan K, Mukherjee A (2003) Modeling aggregate air-travel itinerary shares: logit model development at a major U.S. airline. *J Air Transp Manag* 9(6):361–369
- Data Base Products Inc. (2010) Superset, Dallas, TX, USA
- Erdmann A, Kiahaschemi M, Noltemeier A, Schrader R (1997) Fleet assignment with respect to itineraries. Presentation at mathematical programming conference, Lausanne
- Ervin F, Lonsdale R, Parker R, Zhang J (2005) The effect of alliances in scheduling and strategic planning. Paper presented at the AGIFORS scheduling and strategic planning conference, Seattle
- Farkas A (1996) The Influence of Network Effects and Yield Management on Airline Assignment Decisions. Ph.D. Dissertation, Department of Civil and Environmental Engineering, Massachusetts Institute of Technology, Cambridge, MA
- Gao G, Johnson E, Smith B (2009) Integrated airline fleet and crew robust planning. *Transp Sci* 43(1):2–16
- Garrow LA (2010) Discrete choice modelling and air travel demand: theory and applications. Ashgate, Aldershot
- Goplan R, Talluri K (1998) The aircraft maintenance routing problem. *Oper Res* 46(2):260–271
- Gramming J, Hujer R, Scheidler M (2005) Discrete choice modeling in airline network management. *J Appl Econom* 20:467–486
- Green T (2002) Evaluating airline schedules for improved operational dependability. In: *Proceedings of the AGIFORS annual symposium*. Honolulu, Hawaii
- Gronkvist M (2005) The tail assignment problem. Thesis submitted to the Department of Computer Science and Engineering, Chalmers University of Technology and Goteborg University, Goteborg, Sweden

- Hane C, Barnhart C, Johnson EL, Marsten R, Nemhauser GL, Sigismendi G (1995) The fleet assignment problem: solving a large-scale integer program. *Math Program* 70:211–232
- Jacobs TL (2005) Quantifying market performance risks: a probabilistic approach. In: *Proceedings of the AGIFORS scheduling and strategic planning group*, Seattle, Washington
- Jacobs TL, Johnson EL, Smith BC (1999) O&D FAM: incorporating passenger flows into the fleet assignment process. In: *AGIFORS symposium proceedings 39*, New Orleans
- Jacobs TL, Hunt E, Korol M (2001) Scheduling and revenue management process integration: benefits and hurdles. In: *Proceedings of the AGIFORS revenue management and distribution study group*. Bangkok, Thailand
- Jacobs TL, Ratliff RM, Smith BC (2000) Soaring with synchronized systems: coordinated airline scheduling, yield management and pricing decisions can make airline revenue rake off. *ORMS Today* 27:4
- Jacobs TL, Smith BC, Johnson EL (2008) Incorporating network flow effects into the fleet assignment process. *Transp Sci* 42(4):514–529
- Jacobs TL, Smith BC, Johnson E (2000) System and method for incorporating origination and destination effects into a vehicle assignment process. US patent no. 6,076,067, United States Patent Office, Washington D.C. (Commercially known as O&D FAMTM)
- Kliwer G (2000) Cooperative approaches for market modeling and fleet assignment. Presentation to ISMP 2000. University of Paderborn, Germany
- Knicker TS (1998) Itinerary-based fleet assignment, Ph.D. Dissertation, Massachusetts Institute of Technology
- Koppelman FS, Coldren GC, Parker RA (2008) Schedule delay impacts on air-travel itinerary demand. *Transp Res Part B* 42(3):263–273
- Lacasse-Guay E, Desaulniers G, Soumis F (2010) Aircraft routing under different business processes. *J Air Transp Manag* 16(5):258–263
- Lee LH, Huang HC, Lee C, Chew EP, Jaruphongsa W, Yong YY, Liang Z, Leong CH, Tan YP, Namburi K, Johnson EL, Banks J (2003) Simulation of airports/aviation systems: discrete event simulation model for airline operations: SIMAIR. In: *Proceedings of the 35th conference on Winter simulation: driving innovation*, New Orleans, Louisiana
- Lohatepanont M (2001) Airline fleet assignment and schedule design: integrated models and algorithms, Ph.D. Dissertation, Massachusetts Institute of Technology
- McFadden D (1974) Conditional logit analysis of qualitative choice behavior. In: Zarembka P (ed) *Frontiers in econometrics*. Academic Press, New York, pp 105–142
- Nako SM (1992) Frequent Flyer programs and business travellers: an empirical investigation. *Logist Transp Rev* 28(4):395–414
- Nemhauser GL, Wolsey LA (1988) *Integer and combinatorial optimization*. Wiley, New York
- OAG/OAG Worldwide Limited (2010) *Official Airline Guide*, Bedfordshire, LU5 4HB, UK
- Parker R (2008) The connectedness of the world's airline network: assessing the error introduced when isolating sub-networks of the world's airline networks. In: *Proceedings of the 48th AGIFORS annual symposium*, Montreal, Canada
- Parker RG, Rardin RL (1988) *Discrete optimization*. Academic Press, Boston
- Parker R, Lonsdale R, Glans T, Zhang Z (2005) The na universal market simulator. Presentation to the advanced market techniques forum of the american marketing association, Coeur d'Alene, Idaho
- Prousaloglou K, Koppelman FS (1999) The choice of air carrier, flight and fare class. *J Air Transp Manag* 5(4):193–201
- Qi X, Yang J, Yu G (2004) Chapter 50—Scheduling problems in the airline industry. In: Leung JY-T (ed) *Handbook of scheduling algorithms, models and performance analysis*. Chapman and Hall/CRC, New York
- Ratliff R (2008) Multi-flight demand untruncation: simulation results. In: *Proceedings of the AGIFORS revenue management and distribution study group meeting*. Tahiti, French Polynesia
- Reeves CR (1993) *Modern heuristic techniques for combinatorial problems*. Wiley, New York

- Rosenberger JM, Schaefer AJ, Goldsman D, Johnson EL, Kleywegt AJ, Nemhauser GL (2002) A stochastic model of airline operations. *Transp Sci* 36:357–377
- Ruther S (2010) A multi-commodity flow formulation for the integrated aircraft routing, crew pairing, and tail assignment problem. In: Proceedings of the 45th annual conference of the ORSNZ, November 2010
- Sandhu R, Klabjan D (2007) Integrated airline fleetting and crew pairing decisions. *Oper Res* 55:430–438
- Smith BC, Jacobs TL (1997) Airline planning and marketing in a changing environment. *J Infrastruct Syst* 3(1):1–3
- Smith BC, Johnson EL (2006) Robust airline fleet assignment: imposing station purity using station decomposition. *Transp Sci* 40(4):497–516
- Subramanian R, Scheff RP Jr, Quillinan JD, Wiper DS, Marsten RE (1994) Coldstart: fleet assignment at delta air lines. *Interfaces* 24(1):104–120
- Suzuki Y, Tyworth JE, Novack RA (2001) Airline market share and customer service quality: a reference-dependent model. *Transp Res Part A* 35(9):773–788
- Swan WM (1983) Traffic losses at high load factors. In: AGIFORS symposium proceedings 23, Olive Branch, Mississippi
- Talluri K (1998) The four-day aircraft maintenance routing problem. *Transp Sci* 32(1):43–53
- Transportation Research Board (2002) Transportation research E-circular E-C040, aviation demand forecasting: a survey of methodologies. Transportation Research Board, Washington D.C
- Walker A (2001) Macro-level scheduling and pricing using a genetic algorithm. In: Proceedings of the 41st AGIFORS annual symposium. Sydney, Australia
- Wong D, Jacobs TL (2009) A probabilistic approach to quantifying network & market performance risk. In: Proceedings of the 49th AGIFORS annual symposium. Atlanta, Georgia
- Zeid MA, Rossi TF, Gardner B (2006) Modeling time of day choice in the context of tour and activity based models. Paper presented at the 2006 annual meeting of the transportation research board, Washington D.C
- Zhe L, Chaovalitwongse WA, Huang HC, Johnson EL (2011) Network model for aircraft routing problem. *Transp Sci* 45(1):109–120

Quantitative Problem Solving Methods in the Airline
Industry

A Modeling Methodology Handbook

Barnhart, C.; Smith, B. (Eds.)

2012, X, 462 p., Hardcover

ISBN: 978-1-4614-1607-4