Airline Revenue Management:
An Overview of OR Techniques 1982-2001

Kevin Pak*
Nanda Piersma†

January, 2002

Econometric Institute Report EI 2002-03

Abstract

With the increasing interest in decision support systems and the continuous advance of computer science, revenue management is a discipline which has received a great deal of interest in recent years. Although revenue management has seen many new applications throughout the years, the main focus of research continues to be the airline industry. Ever since Littlewood (1972) first proposed a solution method for the airline revenue management problem, a variety of solution methods have been introduced. In this paper we will give an overview of the solution methods presented throughout the literature.

Keywords: Revenue Management, Seat Inventory Control, OR techniques

* Erasmus Research Institute of Management and Econometric Institute, Erasmus University Rotterdam, The Netherlands
† Econometric Institute, Erasmus University Rotterdam, The Netherlands

1. Introduction

1.1. Revenue Management

Companies selling perishable goods or services often face the problem of selling a fixed capacity of a product over a finite horizon. If the market is characterized by customers willing to pay different prices for the product, it is often possible to target different customer segments by the use of product differentiation. This creates the opportunity to sell the product to different customer segments for different prices, e.g. charging different prices at different points in time or offering a higher service level for a higher price. In order to do so, decisions will have to be made about the prices to charge and the number of products to reserve for each customer segment. Making this kind of decisions is the topic of revenue management.

Revenue management can be defined as the art of maximizing profit generated from a limited capacity of a product over a finite horizon by selling each product to the right customer at the right time for the right price. It encompasses practices such as price-discrimination and turning down customers in anticipation of other, more profitable customers. Revenue management originates from the airline industry, where deregulation of the fares in the 1970's led to heavy competition and the opportunities for revenue management schemes were acknowledged in an early stage. The airline revenue management problem has received a lot of attention throughout the years and continues to be of interest to this day. Other applications of revenue management can be found in the hotel, car rental, railway and cruise-line industries among others. The possible applications of revenue management go beyond the tourist industries, though. The energy and television broadcast industries have been mentioned as possible applications and it has been argued that the concept of revenue management can even be applied to fast moving consumer goods in supermarkets.

1.2. Airline Revenue Management

An airline, typically, offers tickets for many origin-destination itineraries in various fare classes. These fare classes not only include business and economy class, which are settled in separate parts of the plane, but also include fare classes for which the difference in fares is explained by different conditions for e.g. cancellation options or overnight stay arrangements. Therefore the seats on a flight are products which can be offered to different customer segments for different prices. Since the tickets for a flight have to be sold before the plane takes off, the product is perishable and revenue management can be applied.

At the heart of airline revenue management lies the seat inventory control problem. This problem concerns the allocation of the finite seat inventory to the demand that occurs over time before the flight is scheduled to depart. The objective is to find the right combination of passengers on the flights such that revenues are maximized. The optimal allocation of the seat inventory then has to be translated into a booking control policy, which determines whether or not to accept a booking request when it arrives. It is possible that at a certain point in time it is more profitable to reject a booking request in order to be able to accept a booking request of another passenger at a later point in time.

Other important topics that have received attention in the revenue management literature are demand forecasting, overbooking and pricing. Demand forecasting is of critical importance in airline revenue management because booking control policies make use of demand forecasts to determine the optimal booking control strategy. If an airline uses poor demand estimates, this will result in a booking control strategy which performs badly. Airlines often have to cope with no-shows, cancellations and denied boardings. Therefore, in order to prevent a flight from taking off with vacant seats, airlines tend to overbook a flight. This means that the airline books more passengers on a flight than the capacity of the plane allows. The level of overbooking for each type of passenger has been the topic of research for many years. Pricing is obviously very important for the revenues of an airline company. In fact, price differentiation is the starting point of the revenue management concept. Demand forecasting, overbooking and pricing are, however, topics beyond the scope of this paper. For an overview of the literature on these three topics we refer to McGill and Van Ryzin (1999).

2. Seat Inventory Control

The seat inventory control problem in airline revenue management concerns the allocation of the finite seat inventory to the demand that occurs over time. In order to decide whether or not to accept a booking request, the opportunity costs of losing the seats taken up by the booking have to be evaluated and compared to the revenue generated by accepting the booking request. Solution methods for the seat inventory control problem are concerned with approximating these opportunity costs and incorporating them in a booking control policy such that expected future revenues are maximized.

Solution methods for the seat inventory control problem should account for a number of things. The stochastic nature of demand is one of them. Also, a booking request that creates the highest possible revenue for the airline should never be rejected whenever a seat is available, not even when the number of seats appointed to this type of passenger by the booking control policy has been reached. In fact, any passenger should be allowed to tap into the capacity reserved for any other lower valued type of passenger. This is the concept of nesting and should be incorporated into the booking control policy. Further, we make the distinction between single leg and network seat inventory control and static and dynamic solution methods.

With single leg seat inventory control, every flight leg is optimized separately. Network seat inventory control is aimed at optimizing the complete network of flight legs offered by the airline simultaneously. Consider a passenger travelling from A to C through B. That is, travelling from A to C using flight legs from A to B and from B to C. If the single leg approach is used, this passenger can be rejected on one of the flight legs because another passenger is willing to pay a higher fare on this flight leg. But by rejecting this demand, the airline loses an opportunity to create revenue for the combination of the two flight legs. If the other flight leg does not get full, it could have been more profitable to accept the passenger to create revenue for both flight legs. Hence, only the network approach takes into account the overall revenue that the passenger creates from its origin to its final destination.

The distinction between static and dynamic so lution methods is a second partitioning that can be considered. Static solution methods generate an optimal allocation of the seats at a certain point in time, typically the beginning of the booking period, based on a demand forecast at that point in time. The actual booking requests do, however, not arrive at one point in time but occur gradually over the booking period. Therefore, a better solution method would be one that monitors the actual demand and adjusts the booking control policy to it. This would be a dynamic solution method.

In Section 3 we discus the single leg solution methods and in section 4 the network solution methods to the seat inventory control problem. The solution methods may vary with the set of assumptions made in each research, e.g. taking nesting or network-effects into account or not. However, there are also some assumptions that all of the researches discussed in this paper make use of. These assumptions are:

- no cancellations or no-shows
- independent demand between the booking classes
- no demand recapturing
- no batch booking

The first assumption simply states that no attention will go out to overbooking. Usually the seat inventory control problem and overbooking are considered separately, although integration of the two problems would be preferred and has been given attention also. A consequence of the second assumption is that no information on the actual demand process of one fare can be derived from the actual demand process of another fare. We speak of demand recapturing when a low fare booking request is turned into a higher fare booking request when the low fare class is not available. This can occur when the products are not sufficiently differentiated. The assumption that there is no demand recapturing implies that every customer has got a strict preference for a certain fare class and that a denied request is lost forever. The last assumption is that there are no batch bookings, which justifies looking at one booking request at a time. Relaxation of these assumptions has been given attention. However, in order to give a good impression of what is considered as the general seat inventory control problem and its basic solution methods, we will not discuss this here.

Finally, we would like to mention that the seat inventory control problem can also be seen as a pricing problem. When the fare classes are well differentiated, they are separate products. A pricing scheme can then be constructed for each fare class and closing a fare class for future booking requests can be done artificially by setting the price sufficiently high. In our opinion, however, the decision whether to close a fare class or not, can be represented by more straightforward formulations than that of a pricing problem. Whenever the fare classes are not sufficiently differentiated, the fare classes can be seen as different prices for the same product. Then a formulation of the problem as a pricing problem is evident. In this paper, we will, however, not consider this situation. Applications of pricing techniques to airline revenue management can be found in Chatwin (2000), Feng and Gallego (1995, 2000), Feng and Xiao (2000a, 2000b), Gallego and van Ryzin (1994, 1997), Kleywegt (2001), You (1999) and Zhao and Zheng (2000) among others.

3. Single Leg Seat Inventory Control

In single leg seat inventory control, booking control policies for the various flight legs are made independent of one another. There are two categories of single leg solution methods; static and dynamic solution methods. In addition to the assumptions given in the previous section, static single leg solution methods make use of the extra assumption that booking requests come in sequentially in order of increasing fare level, i.e. low fare booking requests come in before high fare booking requests. This means that the booking period can be divided into time-periods for which all booking requests belong to the same fare class. In this case, booking control policies can be based on the total demand for each fare class and do not explicitly have to consider the actual arrival process. Brumelle and McGill (1993) show that under this assumption a static solution method that limits the number of booking requests to accept for each fare class is optimal as long as no change in the probability distributions of demand is foreseen. Dynamic solution methods do not assume a specific arrival order of the booking requests. In this case, a booking control policy based on the total demand for each fare class is no longer optimal, and dynamic programming techniques are needed. In Section 3.1 we discuss the static solution methods and in Section 3.2 the dynamic solution methods.

3.1. Static Solution Methods

Littlewood (1972) was the first to propose a solution method for the seat inventory control problem for a single leg flight with two fare classes. The idea of his scheme is to equate the marginal revenues in each of the two fare classes. He suggests closing down the low fare class when the certain revenue from selling another low fare seat is exceeded by the expected revenue of selling the same seat at the higher fare. That is, low fare booking requests should be accepted as long as

$$f_2 \geq f_1 \Pr(D_1 > p_1) \tag{3.1}$$

where f_1 and f_2 are the high and low fare levels respectively, D_1 denotes the demand for the high fare class, p_1 is the number of seats to protect for the high fare class and $\Pr(D_1 > p_1)$ is the probability of selling all protected seats to high fare passengers. The smallest value of p_1 that satisfies the above condition is the number of seats to protect for the high fare class, and is known as the protection level of the high fare class. The concept of determining a protection level for the high fare class can also be seen as setting a booking limit, a maximum number of bookings, for the lower fare class. Both concepts restrict the number of bookings for the low fare class in order to accept bookings for the high fare class.

Belobaba (1987) extends Littlewood's rule to multiple nested fare classes and introduces the term expected marginal seat revenue (EMSR) for the general approach. His method is known as the EMSRa method and produces nested protection levels, i.e. they are defined as the number of seats protected for the fare class and all higher classes. The EMSRa method does, however, not yield optimal booking limits when more than two fare classes are considered.

Optimal policies for more than two classes have been presented independently by Curry (1990), Brumelle and McGill (1993) and Wollmer (1992). Curry uses continuous demand distributions and Wollmer uses discrete demand distributions. The approach Brumelle and McGill propose, is based on subdifferential optimization and admits either discrete or continuous demand distributions. They show that an optimal set of nested protection levels, $p_1, p_2, ..., p_{k-1}$, where the fare classes are indexed from high to low, must satisfy the conditions:

$$\delta_{+} ER_i(p_i) \leq f_{i+1} \leq \delta_{-} ER_i(p_i) \qquad \text{for each } i = 1, 2, ..., k\text{-}1 \qquad (3.2)$$

where $ER_i(p_i)$ is the expected revenue from the i highest fare classes when p_i seats are protected for those classes and δ_+ and δ_- are the right and left derivatives with respect to p_i respectively. These conditions express that a change in p_i away from the optimal level in either direction will produce a smaller increase in the expected revenue than an immediate increase of f_{i+1}. The same conditions apply for discrete and continuous demand distributions. Notice, that it is only necessary to set k-1 protection levels when there are k fare classes on the flight leg, because no seats will have to be protected for the lowest fare class. Brumelle and McGill show that under certain continuity conditions the conditions for the optimal nested protection levels reduce to the following set of probability statements:

$$\begin{aligned}
f_2 &= f_1 \Pr(D_1 > p_1) \qquad (3.3)\\
f_3 &= f_1 \Pr(D_1 > p_1 \cap D_1 + D_2 > p_2)\\
&\ldots\\
f_k &= f_1 \Pr(D_1 > p_1 \cap D_1 + D_2 > p_2 \cap \ldots \cap D_1 + D_2 + \ldots + D_{k-1} > p_{k-1})
\end{aligned}$$

These statements have a simple and intuitive interpretation, much like Littlewood's rule. Just like Littlewood's rule and the EMSRa method, this method is based on the idea of equating the marginal revenues in the various fare classes and therefore belongs to the class of EMSR methods. The method is called the EMSRb method. Robinson (1995) finds the optimality conditions when the assumption of a sequential arrival order with monotonically increasing fares is relaxed into a sequential arrival order with an arbitrary fare order. Furthermore, Curry (1990) provides an approach to apply his method to origin-destination itineraries instead of single flight legs, when the capacities are not shared among different origin-destinations.

Van Ryzin and McGill (2000) introduce a simple adaptive approach for finding protection levels for multiple nested fare classes, which has the distinctive advantage that it does not need any demand forecasting. Instead, the method uses historical observations to guide adjustments of the protection levels. They suggest adjusting the protection level p_i upwards after each flight if all the fare classes i and higher reached their protection levels, and downwards if this has not occurred. They prove that under reasonable regularity conditions, the algorithm converges to the optimal nested protection levels. This scheme of continuously adjusting the protection levels has the advantage that it does not need any demand forecasting and therefore is a way to get around all the difficulties involving this practice. However the updating scheme does need a sufficiently large sequence of flights to converge to a good set of protection levels. In practice, such a start-up period can not always be granted when there are profits to be made.

The solution methods in this paragraph are all static. This class of solution methods is optimal under the sequential arrival assumption as long as no change in the probability distributions of the demand is foreseen. However, information on the actual demand process can reduce the uncertainty associated with the estimates of demand. Hence, repetitive use of a static method over the booking period based on the most recent demand and capacity information, is the general way to proceed.

3.2. Dynamic Solution Methods

Dynamic solution methods for the seat inventory control problem do not determine a booking control policy at the start of the booking period as the static solution methods do. Instead, they monitor the state of the booking process over time and decide on acceptance of a particular booking request when it arrives, based on the state of the booking process at that point in time.

Lee and Hersh (1993) consider a discrete-time dynamic programming model, where demand for each fare class is modeled by a nonhomogeneous Poisson process. Using a Poisson process gives rise to the use of a Markov decision model in such a way that, at any given time t, the booking requests before time t do not affect the decision to be made at time t except in the form of less available capacity. The states of the Markov decision model are only dependent on the time until the departure of the flight and on the remaining capacity. The booking period is divided into a number of decision periods. These decision periods are sufficiently small such that not more than one booking request arrives within such a period. The state of the process changes every time a decision period elapses or the available capacity changes. If $U(c,t)$ is the optimal total expected revenue that can be generated given a remaining capacity of c seats and with t remaining decision periods before the departure of the flight, then a booking request of class i is accepted if, and only if:

$$f_i \geq U(c,t-1) - U(c-1,t-1) \qquad \text{for each } i = 1, 2, ..., k, \qquad (3.4)$$
$$c = C, C\text{-}1, ..., 1, \; t = T, T\text{-}1, ..., 1$$

where C is the total seat capacity and T is the total number of decision periods. This decision rule says that a booking request is only accepted if its fare exceeds the opportunity costs of the seat, defined here by the expected marginal value of the seat at time t. Lee and Hersh provide a recursive function for the total expected revenue and show that solving the model under the decision rule given by (3.4) results into a booking policy that can be expressed as a set of critical values for either the remaining capacity or the time until departure. For each fare class the critical values provide either an optimal capacity level for which booking requests are no longer accepted in a given decision period, or an optimal decision period after which booking requests are no longer accepted for a given capacity level. The critical values are monotone over the fare classes. Lee and Hersh also provide an extension to their model to incorporate batch arrivals.

Kleywegt and Papastavrou (1998) demonstrate that the problem can also be formulated as a dynamic and stochastic knapsack problem (DSKP). Their work is aimed at a broader class of problems than only the single leg seat inventory control problem considered here, and includes the possibility of stopping the process before time 0 with a given terminal value for the remaining capacity, waiting costs for capacity unused and a penalty for rejecting an item. Their model is a continuous-time model, but they do, however, only consider homogeneous arrival processes for the booking requests. In a recent paper Kleywegt and Papastavrou (2001) extend their model to allow for batch arrivals.

Subramanian et al. (1999) extend the model proposed by Lee and Hersh to incorporate cancellations, no-shows and overbooking. They also consider a continuous-time arrival process as a limit to the discrete-time model by increasing the number of decision periods. Liang (1999) reformulates and solves the Lee and Hersh model in continuous-time. Van Slyke and Young (2000) also obtain continuous-time versions of Lee and Hersh' results. They do this by simplifying the DSKP model to the more standard single leg seat inventory control problem and extending it for nonhomogeneous arrival processes. They also allow for batch arrivals. Lautenbacher and Stidham (1999) link the dynamic and static approaches. They demonstrate that a common Markov decision process underlies both approaches and formulate an omnibus model which encompasses the static and dynamic models as special cases.

4. Network Seat Inventory Control

In network seat inventory control, the complete network of flights offered by the airline is optimized simultaneously. One way to do this, is to distribute the revenue of an origin-destination itinerary over its legs, which is called prorating, and apply single leg seat inventory control to the individual legs. Williamson (1992) investigates different prorating strategies, such as prorating based on mileage and on the ratio of the local fare levels. This approach provides a heuristic to extend the existing single leg solution methods to a network setting. However, only a mathematical programming formulation of the problem can be capable of fully capturing the combinatorial aspects of the network. In order to obtain the mathematical programming formulation for capturing these combinatorial aspects, denote an origin-destination and fare class combination by ODF. Let X_{ODF} denote the number of seats reserved for an ODF, D_{ODF} the demand for an ODF, and f_{ODF} the fare level for an ODF. Further, let l denote a single flight leg, C_l the seat capacity for a leg, and S_l the set of all ODF combinations available on a leg. The problem can then be formulated as follows:

$$
\begin{aligned}
&\text{maximize} && E\left(\textstyle\sum_{ODF} f_{ODF} \min\{X_{ODF}, D_{ODF}\}\right) && && (4.1)\\
&\text{subject to} && \textstyle\sum_{ODF \in S_l} X_{ODF} \le C_l && \text{for each } l\\
& && X_{ODF} \ge 0 \text{ integer} && \text{for each } ODF
\end{aligned}
$$

The objective is to find the seat allocation that maximizes the total expected revenue of the network and satisfies the capacity constraints on the various flight legs. The objective function depends on the distributions of demand and generally is not linear, continuous or in any other way regular. Therefore, relaxations of this formulation have been suggested for use in practice.

4.1 Mathematical Programming

The first full network formulation of the seat inventory control problem is proposed by Glover et al. (1982). They formulate the problem as a minimum cost network flow problem with one set of arcs corresponding to the flight legs and another set corresponding to the ODF combinations. The method is aimed at finding the flow on each arc in the network that maximizes revenue, without violating the capacity constraints on the legs and upperbounds posed by the demand forecasts for the ODF combinations. A drawback of the network flow formulation is that it can not always discriminate between the routes chosen from an origin to a destination. Therefore, this formulation only holds when passengers are path-indifferent. The advantage of the formulation is that is it easy to solve and can be re-optimized very fast.

A formulation of the problem that is able to distinguish between the different routes from an origin to a destination, is given by the integer programming model underlying the network flow formulation:

$$\begin{array}{lll} \text{maximize} & \sum_{ODF} f_{ODF} X_{ODF} & \\ \text{subject to} & \sum_{ODF \in S_l} X_{ODF} \leq C_l & \text{for each } l \\ & X_{ODF} \leq ED_{ODF} & \text{for each } ODF \\ & X_{ODF} \geq 0 \ \text{ integer} & \text{for each } ODF \end{array} \tag{4.2}$$

In this model ED_{ODF} denotes the expected demand for an ODF. It is easy to see that this is the model obtained from model (4.1) if the stochastic demand for each ODF is replaced by its expected value. The demand for an ODF is treated as if it takes on a known value, e.g. as if it is deterministic, and no information on the demand distributions is taken into account. Accordingly, the model produces the optimal seat allocation if the expected demands correspond perfectly with the actual demands. It is common practice to solve the LP relaxation of the model rather than the integer programming problem, since an integer programming problem is usually very hard to solve. The LP relaxation of the model is known as the deterministic linear programming (DLP) model. A booking control policy based on the DLP model can be constructed by setting booking limits for each ODF equal to the number of seats reserved for the ODF in the optimal solution of the model. Such a booking control policy is a static method and, just as with the single leg methods discussed in the previous section, the general way to proceed is to use the model repeatedly over the booking period based on the most recent demand and capacity information.

The DLP method is a deterministic method and will never reserve more seats for a higher fare class than the airline expects to sell on average. In order to determine whether reserving more seats for more profitable ODF combinations can be rewarding, it is necessary to incorporate the stochastic nature of demand in the model. Wollmer (1986) develops a model which incorporates probabilistic demand into a network setting.

$$\text{maximize} \quad \sum_{ODF} \sum_{i} f_{ODF} \Pr(D_{ODF} \geq i) X_{ODF}(i) \tag{4.3}$$

$$\text{subject to} \quad \sum_{ODF \in S_l} \sum_{i} X_{ODF}(i) \leq C_l \qquad \text{for each } l$$

$$X_{ODF}(i) \in \{0,1\} \qquad \text{for each } ODF,$$

$$i = 1, 2, ..., \max_l\{C_l : ODF \in S_l\}$$

In this model the decision variables $X_{ODF}(i)$ take on the value 1 when i seats or more are reserved for the ODF, and 0 otherwise. The coefficient of each $X_{ODF}(i)$ in the objective function represents the expected marginal revenue of allocating an additional i^{th} seat to the ODF. The model is called the expected marginal revenue (EMR) model. A drawback of this model is the large amount of decision variables, which makes the model impractical in use.

De Boer et al. (1999) introduce a model which is an extension of the EMR model. It incorporates the stochastic nature of demand when demand for each ODF can take on only a limited number of discrete values $\{d_{ODF}(1) < d_{ODF}(2) < ... < d_{ODF}(N_{ODF})\}$.

maximize	$\sum_{ODF}\sum_i f_{ODF}\ \Pr(D_{ODF} \geq d_{ODF}(i))X_{ODF}(i)$		(4.4)
subject to	$\sum_{ODF\in S_l}\sum_i X_{ODF}(i) \leq C_l$	for each l	
	$X_{ODF}(1) \leq d_{ODF}(1)$	for each ODF	
	$X_{ODF}(i) \leq d_{ODF}(i) - d_{ODF}(i-1)$	for each ODF, $i = 2, 3, ..., N_{ODF}$	
	$X_{ODF}(i) \geq 0$ integer	for each ODF, $i = 1, 2, ..., N_{ODF}$	

The decision variables $X_{ODF}(i)$ each accommodate for the part of the demand D_{ODF} that falls in the interval ($d_{ODF}(i\text{-}1)$, $d_{ODF}(i)$]. Summing the decision variables $X_{ODF}(i)$ over all i for an ODF, gives the total number of seats reserved for the ODF which can be interpreted as a booking limit. The LP relaxation of this model is called the stochastic linear programming (SLP) model. The EMR model is a special case of the SLP model that can be obtained by letting $d_{ODF}(1) = 1$ and $d_{ODF}(i)\text{-}d_{ODF}(i\text{-}1) = 1$ for all $i = 2, 3, ..., \max_l\{C_l{:}ODF\in S_l\}$. But the SLP formulation of the problem is more flexible because it allows a reduction of the number of decision variables by choosing a limited amount of demand scenarios. If only the expected demand is considered as a possible scenario, the SLP model reduces to the DLP model. In fact, the DLP and EMR models can be seen as the two extremes that can be obtained from the SLP model. The first by considering only one demand scenario, the latter by considering all possible scenarios.

The mathematical programming models discussed in this section are very well capable of capturing the combinatorial aspects of the problem. The booking control policies derived from the models are, however, static and non-nested. In the following sections we will discuss techniques to augment the mathematical programming models for nesting. Dynamic solution methods are discussed in Section 4.2.

4.1.1. Nesting

Nesting is an important aspect of the seat allocation problem and should be taken into account. How to determine a nesting order of the ODF combinations is not trivial in a network setting. The nesting order should be based on the contribution of the ODF combinations to the network revenue. Ordering by fare class does not take into account the level of the fare, and ordering by fare level does not account for the load factors of the flight legs. Williamson (1992) suggests nesting the ODF combinations by the incremental revenue that is generated if an additional seat is made available for the ODF while everything else remains unchanged. For the DLP model, she approximates this by the dual price of the corresponding demand constraint. In this particular model, this corresponds to the incremental revenue obtained from increasing the mean demand for the ODF by one. A stochastic model typically does not have demand constraints, but the incremental revenue obtained from increasing the mean demand can still be used. An approximation can be obtained by re-optimizing the model with the mean demand increased by one and comparing the new objective value with the original objective value. This does, however, require a re-optimization of the model.

After determining a nesting order on each flight leg, a nested booking control policy can be constructed. Let $H_{ODF,l}$ be the set of ODF combinations that have higher rank than *ODF* on flight leg l. Then nested booking limits for an ODF on a flight leg l are given by:

$$b_{ODF,l} = C_l - \sum_{ODF^* \in H_{ODF,l}} X_{ODF^*} \tag{4.5}$$

This illustrates that nested booking limits are obtained from non-nested booking limits by allowing ODF combinations to make use of all seats on the flight leg except for the seats reserved for higher ranked ODF combinations.

De Boer et al. (1999) stick to Williamson’s idea of using the net contribution to network revenue of the ODF combinations to determine a nesting. However, they use a different approach to approximate this. They approximate the opportunity costs of an ODF combination by the sum of the dual prices of the capacity constraints of the legs the ODF uses. An approximation of the net contribution to network revenue is then obtained by subtracting this from the fare level. Thus, a nesting order is based on:

$$\overline{f}_{ODF} = f_{ODF} - \sum_{ODF \in S_l} p_l \tag{4.6}$$

where p_l denotes the dual price of the capacity constraint for flight leg l. For the DLP method this nesting method is equivalent to Williamson's approach. The advantage of this method over Williamson's, is that it can be applied for a stochastic model without re-optimizing the model.

4.1.2. Bid-Prices

A booking control policy that incorporates nesting in a natural way, is setting bid-prices. In this procedure, a bid-price is set for each leg in the network reflecting the opportunity costs of reducing the capacity of the leg with one seat. A booking request is accepted only if its fare exceeds the sum of the bid-prices of the legs it uses. The opportunity costs of selling a seat on a leg can be approximated by the dual price of the capacity constraint of the leg in a mathematical programming model. After obtaining the dual prices of the capacity constraints by the use of such a model, the rule is to accept a booking request for an ODF if:

$$f_{ODF} > \sum_{ODF \in S_l} p_l \qquad (4.7)$$

Notice that this measure is equivalent to the approximation de Boer et al. (1999) use for the opportunity costs of an ODF and directly links the revenue gain from accepting a booking request to the opportunity costs of the ODF. A disadvantage of bid-price control is that there is no limit to the number of bookings for an ODF once it is open for bookings, i.e. once its fare exceeds the opportunity costs. This can lead to flights filling up with passengers that only marginally contribute to network revenue. Frequently adjusting the bid-prices based on the most recent demand and capacity information is necessary to prevent this from happening.

Williamson (1992) investigates using the DLP model for constructing bid-prices. This method to construct bid-prices does not take into account the stochastic nature of demand. Talluri and van Ryzin (1999) analyze a randomized version of the DLP method for computing bid-prices. The idea is to incorporate more stochastic information by replacing the expected demand by the random vector itself. They simulate a sequence of n demand realizations and for each realization determine the optimal seat allocation. This can be done by applying the DLP model with the realization of the demand taking the place of the expected demand as the upperbound for the number of bookings for each ODF. The n optimal seat allocations provide n sets of dual prices. The bid-price for a leg is simply defined as the average over the n dual prices for the flight leg. This method is known as the randomized linear programming (RLP) method. De Boer et al. (1999) construct bid-prices on their SLP model.

It should be noted that both the nested booking limits and the bid-price procedures are heuristics to convert a non-nested solution from one of the mathematical programming models into a nested booking control policy by allowing ODF combinations to make use of all seats reserved for the lower valued ODF combinations. Allowing this, reduces the necessity to reserve seats for the ODF in the model. Therefore, the solution of the model is no longer optimal. To obtain an optimal booking control strategy that accounts for nesting, the nesting and allocation decisions should be integrated. No mathematical programming model is capable of doing this. A heuristic that does integrate the nesting and allocation decisions is discussed in the next section.

4.2. Simulation Approach

In a recent study, Bertsimas and de Boer (2000) introduce a simulation based solution method for the network seat inventory control problem. They define the expected revenue function as a function of the booking limits and their aim is to find those booking limits that optimize the function. The DLP model is used to generate an initial solution which takes the combinatorial aspects of the network into account and by which a nesting order can be determined. After that, the solution is gradually improved to make up for factors such as the stochastic nature of demand and nesting. The search direction is determined by the gradient of the expected revenue function. Because the expected revenue function is not known, it is approximated by means of simulation. The expected revenue generated by a set of booking limits is approximated by the average of the revenues generated by the booking limits when they are applied over a sequence of simulated demand realizations. The gradient of the function is approximated by the change in expected revenue caused by a small deviation in the booking limits.

Bertsimas and de Boer reduce a great deal of the computation time of their method by linking it to ideas from the field of approximate dynamic programming. They devide the booking-period into smaller time-periods and define future revenue as a function of the remaining capacity. A booking control policy for the current time-period can then be obtained by simulating the booking process of the present time-period only. The revenue of each simulation run is defined as the revenue within the present time-period plus the estimated future revenue which depends on the remaining capacity. In order to estimate the future revenue function an Orthogonal Array/Multiple Adaptive Regression Splines method is used as in Chen et al. (1998), which we will discuss in the next section when we present the dynamic solution methods for network seat inventory control.

Bertsimas and de Boer also provide a method to derive bid-prices from their booking limits by use of simulation. The bid-price for each leg is set equal to an approximation of the opportunity costs of reducing the capacity on the leg. They simulate a sequence of demand realizations and for each simulation calculate the revenue resulting from using the booking limits. To obtain an approximation of the opportunity costs, they subtract from this revenue the revenue generated by the same booking limits if the capacity on the leg would have been decreased by one seat. The bid-price is defined as the average of the approximated opportunity costs over the simulations.

4.3. Dynamic Solution Methods

For the simulation based solution method discussed in the previous section, Bertsimas and de Boer (2000) make use of approximate dynamic programming. They devide the booking period into smaller time-periods for which booking control policies are determined. A solution is constructed in each period taking into account the realizations in the previous time-periods and the expectations about the future time-periods. All other network solution methods discussed thus far, are static methods. These methods produce a solution at a given point in time for the complete booking period. This solution is usually adjusted a multitude of times during the booking period by re-optimizing the underlying models. A fully dynamic solution method, however, would be one that adjusts the booking control policy continuously.

Chen et al. (1998) are the first to provide a fully dynamic solution method for the network seat inventory control problem. They formulate a Markov decision model that uses mathematical programming in a dynamic setting. As with the single leg dynamic solution methods, the state space of the Markov decision model is defined by the time until departure and the remaining capacities of the flights. The decision periods are chosen sufficiently small such that not more than one booking request arrives within such a period. Let $V(\boldsymbol{c},t)$ be the optimal total expected revenue that can be generated when $\boldsymbol{c}$ is the vector of remaining capacities on the flight legs and t is the number of decision periods left before departure. Further, let $\boldsymbol{a}_{\boldsymbol{ODF}}$ be the vector that denotes whether a flight leg is used by an ODF or not; i.e. 1 if the ODF traverses the flight leg and 0 otherwise. Then a booking request for an ODF is accepted if, and only if:

$$f_{ODF} \geq V(\boldsymbol{c}, t-1) - V(\boldsymbol{c} - \boldsymbol{a}_{\boldsymbol{ODF}}, t-1) \qquad \text{for each } ODF, \boldsymbol{c}, \; t = T, T\text{-}1, ..., 1 \qquad (4.8)$$

where T is the total number of decision periods. The right-hand side of (4.8) corresponds to the opportunity costs of the seats taken up by the booking request. A booking request is accepted only if its fare exceeds the opportunity costs.

To approximate the opportunity costs, the objective value for a mathematical programming model can be evaluated when the booking request is accepted as well as when the booking request is rejected. Subtracting these objective values gives the opportunity costs based on that particular model. Chen et al. (1998) argue that the opportunity costs are overestimated by the DLP model and underestimated by a non-nested stochastic model they formulate. Based on this idea, they formulate the following algorithm to accept or reject a booking request for an ODF:

1. reject if $f_{ODF} \leq OC_{STOCH}$, otherwise
2. accept if $f_{ODF} \geq OC_{DLP}$, otherwise
3. accept if $f_{ODF} > x$, with x random from the interval $[OC_{STOCH}, OC_{DLP}]$.

where OC_{STOCH} and OC_{DLP} denote the opportunity costs of the ODF as approximated by the stochastic and the DLP model. Evaluating the two models in two different states every time a booking request comes in, obviously requires a lot of computation time. Therefore, Chen et al. propose a method to estimate the value function of a model for each possible state beforehand. They evaluate the model on a carefully selected limited number of points in the state space and use these observations to estimate the value function of the model over the entire state space. The selection of the points is based on an Orthogonal Array method, and Multivariate Adaptive Regression Splines are used to estimate the value function of the model. With an approximation of the value function of each model available at any time, the Markov decision model can be used in a dynamic way.

In a recent paper Bertsimas and Popescu (2001) use the network flow formulation of the problem, proposed by Glover et al. (1982), to approximate the opportunity costs. Because this formulation can be re-optimized very efficiently, a new solution can be constructed every time a booking request comes in. Bertsimas and Popescu overcome the fact that the network flow formulation does not account for the stochastic nature of demand by means of simulation. They simulate a sequence of demand realizations and approximate the opportunity costs by the average of the opportunity costs obtained from the simulations. A drawback of the network flow formulation remains that it only holds when passengers are path-indifferent.

4. Conclusion

In this paper, we make a distinction between single leg and network solution methods for the seat inventory control problem in airline revenue management. Apart from the distinction between static and dynamic solution methods, literature on the single leg approach to the problem is rather harmonious. For both the static and the dynamic approach, a certain amount of consensus has been reached about the general way to proceed. In recent years, literature on single leg solution methods has been aimed mainly at extending the existing models to account for aspects such as overbooking, batch arrivals, less dependence on demand forecasts etc. Literature on the network solution methods is less harmonious. How to account for the combinatorial effects of the network, the stochastic nature of demand and nesting simultaneously, is not trivial. Moreover, the size of the problem prescribes the use of heuristics as opposed to optimal policies, especially if a policy is to be used in a dynamic way. Nevertheless, we think that it is essential to account for the network effects.

5. References

Belobaba, P.P. (1987), Airline Yield Management: An Overview of Seat Inventory Control, Transportation Science, 21, 63-73.

Bertsimas, D. and de Boer, S.V. (2000), Simulation Based Optimization for Airline Network Yield Management, Working Paper, Massachusetts Institute of Technology, Cambridge, MA.

Bertsimas, D. and Popescu, I. (2001), Revenue Management in a Dynamic Network Environment, Working Paper, Massachusetts Institute of Technology, Cambridge, MA.

Brumelle, S.L. and McGill, J.I. (1993), Airline Seat Allocation with Multiple Nested Fare Classes, Operations Research, 41, 127-137.

Chatwin, R.E. (2000), Optimal Dynamic Pricing of Perishable Products with Stochastic Demand and a Finite Set of Prices, European Journal of Operational Research, 125, 149-174.

Chen, V.C.P., Günther, D. and Johnson, E.L. (1998), A Markov Decision Problem Based Approach to the Airline YM Problem, Working Paper, The Logistics Institute, Georgia Institute of Technology, Atlanta, GA.

Curry, R.E. (1990), Optimal Airline Seat Allocation with Fare Classes Nested by Origins and Destinations, Transportation Science, 24, 193-204.

De Boer, S.V., Freling, R. and Piersma, N. (1999), Mathematical Programming for Network Revenue Management Revisited, Working Paper, Econometric Institute, Erasmus University Rotterdam, The Netherlands.

Feng, Y.Y. and Gallego, G. (1995), Optimal Starting Times for End-Of-Season Sales and Optimal Stopping-Times for Promotional Fares, Management Science, 41, 1371-1391.

Feng, Y.Y. and Gallego, G. (2000), Perishable Asset Revenue Management with Markovian Time Dependent Demand Intensities, Management Science, 46, 941-956.

Feng, Y.Y. and Xiao, B. (2000a), A Continuous-Time Yield Management Model with Multiple Prices and Reversible Price Changes, Management Science, 46, 644-657.

Feng, Y.Y. and Xiao, B. (2000b), Optimal Policies of Yield Management with Multiple Predetermined Prices, Operations Research, 48, 332-343.

Gallego, G. and van Ryzin, G.J. (1994), Optimal Dynamic Pricing of Inventories with Stochastic Demand over Finite Horizons, Management Science, 40, 999-1020.

Gallego, G. and van Ryzin, G.J. (1997), A Multiproduct Dynamic Pricing Problem and its Applications to Network Yield Management, Operations Research, 45, 24-41.

Glover, F., Glover, R., Lorenzo, J. and McMillan, C. (1982), The Passenger Mix Problem in the Scheduled Airlines, Interfaces, 12, 73-79.

Kleywegt, A.J. (2001), An Optimal Control Problem of Dynamic Pricing, Working Paper, Georgia Institute of Technology, Atlanta, GA.

Kleywegt, A.J. and Papastavrou, J.D. (1998), The Dynamic and Stochastic Knapsack Problem, Operations Research, 46, 17-35.

Kleywegt, A.J. and Papastavrou, J.D. (2001), The Dynamic and Stochastic Knapsack Problem with Random Sized Items, Operations Research, 49, 26-41.

Lautenbacher, C.J. and Stidham jr., S. (1999), The Underlying Markov Decision Process in the Single-Leg Airline Yield Management Problem, Transportation Science, 33, 136-146.

Lee, T.C. and Hersh, M. (1993), A Model for Dynamic Airline Seat Inventory Control with Multiple Seat Bookings, Transportation Science, 27, 252-265.

Liang, Y. (1999), Solution to the Continuous Time Dynamic Yield Management Model, Transportation Science, 33, 117-123.

Littlewood, K. (1972), Forecasting and Control of Passenger Bookings, in AGIFORS Symposium Proc. 12, Nathanya, Israel.

McGill, J.I. and van Ryzin, G.J. (1999), Revenue Management: Research Overview and Prospects, Transportation Science, 33, 233-256.

Robinson, L.W. (1995), Optimal and Approximate Control Policies for Airline Booking with Sequential Nonmonotonic Fare Classes, Operations Research, 43, 252-263.

Subramanian, J., Stidham jr., S. and Lautenbacher, C.J. (1999), Airline Yield Management with Overbooking, Cancellations, and No-Shows, Transportation Science, 33, 147-167.

Talluri, K.T. and van Ryzin, G.J. (1999), An Analysis of Bid-Price Controls for Network Revenue Management, Management Science, 44, 1577-1593.

Van Ryzin, G.J. and McGill, J.I. (2000), Revenue Management Without Forecasting or Optimization: An Adaptive Algorithm for Determining Airline Seat Protection Levels, Management Science, 46, 760-775.

Van Slyke, R. and Young, Y. (2000), Finite Horizon Stochastic Knapsacks with Applications to Yield Management, Operations Research, 48, 155-172.

Williamson, E.L. (1992), Airline Network Seat Inventory Control: Methodologies and Revenue Impacts, Ph.D. thesis, Flight Transportation Laboratory, Massachusetts Institute of Technology, Cambridge, MA.

Wollmer, R.D. (1986), A Hub-Spoke Seat Management Model, unpublished company report, Douglas Aircraft Company, McDonnell Douglas Corporation, Long Beach, CA.

Wollmer, R.D. (1992), An Airline Seat Management Model for a Single Leg Route when Lower Fare Classes Book First, Operations Research, 40, 26-37.

You, P.S. (1999), Dynamic Pricing in Airline Seat Management for Flights with Multiple Flight Legs, Transportation Science, 33, 192-206.

Zhao, W. and Zheng, Y.S. (2000), Optimal Dynamic Pricing for Perishable assets with nonhomogeneous demand, Management Science, 46, 375-388.