INFORMS Simulation Society Research Workshop 2007

NEW STOCHASTIC APPROXIMATION AND APPROXIMATE DYNAMIC PROGRAMMING METHODS FOR REVENUE MANAGEMENT PROBLEMS

Huseyin Topaloglu and Sumit Kunnumkal

School of Operations Research and Information Engineering
223 Rhodes Hall, Cornell University
Ithaca, NY 14853, USA

1 Introduction

A fundamental problem faced by all airlines is that of deciding whether a request for a particular itinerary should be accepted given that the future itinerary requests are uncertain and the capacities on the flight legs are limited. These types of problems, to which we broadly refer to as revenue management problems, have been studied extensively and obtaining good solutions for them is crucial for the airlines. Nevertheless, despite their long history and practical importance, the solution methods for revenue management problems, especially for the versions that take place over complex flight networks, traditionally do not go beyond simple deterministic linear programs. Our goal here is to present several new solution methods that are motivated by the recent developments in the stochastic approximation and approximate dynamic programming literature.

One class of revenue management problems involve allocating the capacity on a single flight leg. The optimal policy for these problems can easily be obtained through dynamic programming as long as the demand distributions are known. In practice, the demand information is usually censored by how many seats are made available for sale and this makes it very difficult to forecast the demand random variables. We address this difficulty by proposing a stochastic approximation method. Our method only requires observing whether the demand strictly exceeds the number seats that we make available. Therefore, it is applicable when the demand information is censored.

A popular method to obtain good solutions for the revenue management problems that take place over multiple flight legs is to use bid prices. This idea associates a bid price with each flight leg and an itinerary request is accepted only if the revenue from the requested itinerary exceeds the sum of the bid prices associated with the flight legs that are in the requested itinerary. Clearly, a fixed set of bid prices parameterizes a control policy and this suggests using a stochastic approximation method to search for a good set of bid prices. However, the difficulty is that since we cannot sell a fraction of a seat, the problem is inherently discrete and we need to work with an appropriately smoothed version of the problem to obtain convergence.

The revenue management problem that takes place over multiple flight legs can also be viewed as a weakly coupled dynamic program. In these dynamic programs, the evolutions of the different components of the state variable are affected by different types of decisions, but these decisions interact through a set of linking constraints. We exploit this view and relax the constraints that link the decisions for different flight legs by associating dual multipliers with them. As a result, the problem decomposes into a sequence of subproblems, each taking place over a single flight leg, and we obtain very tight approximations to the value functions. As a by product, we obtain bid prices that depend on the capacity remaining on different flight legs.

Since our goal is to present a number of different solution methods, we give pointers to the relevant literature as we describe each solution method.

2 A Stochastic Approximation Method for the Single Leg Problem

We want to use c seats available on a single flight leg to satisfy the demands from n fare classes that arrive sequentially. We index the fare classes such that the demand from fare class 1 arrives first and the demand from fare class n arrives last. If we sell a seat to fare class j, then we generate a revenue of r_j. We employ the standard assumption that the revenues satisfy $0 < r_1 \leq r_2 \leq \ldots \leq r_n$ so that the demands from the cheaper fare classes arrive earlier. We let D_j be the demand from fare class j and assume that $\{D_j : j = 1, \ldots, n\}$ are independent random variables that take integer values.

2.1 Dynamic Programming Formulation

Letting x_j be the remaining capacity just before making the decisions for fare class j and u_j be the number of seats sold to fare class j, the optimal policy can be found by solving the optimality equations

$$v_j(x_j, D_j) = \max_{0 \leq u_j \leq \min\{x_j, D_j\}} r_j u_j + \mathbb{E}\big\{v_{j+1}(x_j - u_j, D_{j+1})\big\}.$$

The constraints above ensure that the number of seats sold do not exceed the remaining capacity and the demand from fare class j. Alternatively, if we let $y_j = x_j - u_j$, then the expression above becomes

$$v_j(x_j, D_j) = \left\{ \max_{[x_j - D_j]^+ \leq y_j \leq x_j} -r_j y_j + \mathbb{E}\big\{v_{j+1}(y_j, D_{j+1})\big\} \right\} + r_j x_j. \quad (1)$$

It is possible to show that $\{v_j(\cdot, D_j) : j = 1, \ldots, n\}$ are piecewise-linear concave functions with points of nondifferentiability being a subset of integers w.p.1; see Brumelle and McGill (1993) and Talluri and van Ryzin (2004). In this case, it is easy to see that the optimal policy is characterized by a set of protection levels $\{y_j^* : j = 1, \ldots, n\}$, where y_j^* can be computed as a maximizer of the function

$$f_j(y_j) = -r_j y_j + \mathbb{E}\big\{v_{j+1}(y_j, D_{j+1})\big\} \quad (2)$$

over the interval $[0, c]$. This is to say that if the remaining capacity just before making the decisions for fare class j is x_j and the demand from fare class j is D_j, then it is optimal to sell $\min\{[x_j - y_j^*]^+, D_j\}$ seats to fare class j. Using the fact that the demands from the cheaper fare classes arrive earlier, it is also possible to show that the optimal protection levels are nested. In other words, there exists a set of optimal protection levels $\{y_j^* : j = 1, \ldots, n\}$ such that we have $y_1^* \geq y_2^* \geq \ldots \geq y_n^*$.

2.2 Stochastic Approximation Method

Our goal is to compute the optimal protection levels through a stochastic approximation method. Unfortunately, the total expected revenue, when viewed as a function of the protection levels, is not concave. Therefore, a naive stochastic approximation method to search for the optimal protection levels leaves us with a nonconvex nonsmooth problem and the standard theory in Kushner and Clark (1978), Ermoliev (1988) and Bertsekas and Tsitsiklis (1996) does not ensure that we get the optimal protection levels.

By (2), we can compute a stochastic subgradient of $f_j(\cdot)$ at y_j through

$$\Delta_j(y_j, D_{j+1}) = -r_j + \dot{v}_{j+1}(y_j, D_{j+1}), \quad (3)$$

where we use $\dot{v}_{j+1}(y_j, D_{j+1})$ to denote a stochastic subgradient of $\mathbb{E}\big\{v_{j+1}(\cdot, D_{j+1})\big\}$ at y_j. In this case, letting $\{y_j^k : j = 1, \ldots, n\}$ be the estimates of the optimal protection levels at iteration k, $\{D_j^k : k = 1, \ldots, n\}$ be the demand random variables at iteration k and $\{\alpha_j^k : j = 1, \ldots, n\}_k$ be a sequence of step size parameters, we can update our estimates of the optimal protection levels by

$$y_j^{k+1} = \min\big\{\big[y_j^k + \alpha_j^k \Delta_j(y_j^k, D_{j+1}^k)\big]^+, c\big\},$$

where the operator $\min\{[\cdot]^+, c\}$ ensures that the estimates of the optimal protection levels always lie in the interval $[0, c]$. However, this approach is clearly not realistic because the computation in (3) requires the knowledge of $\{v_j(\cdot, \cdot) : j = 1, \ldots, n\}$. The method we propose in this section is based on constructing tractable approximations to the stochastic subgradients of $\{f_j(\cdot) : j = 1, \ldots, n\}$.

Since $f_j(\cdot)$ is concave and the optimal protection level y_j^* is a maximizer of this function over the interval $[0, c]$, we can use a standard argument to write (1) as

$$v_j(x_j, D_j) = \begin{cases} r_j D_j + \mathbb{E}\big\{v_{j+1}(x_j - D_j, D_{j+1})\big\} & \text{if } y_j^* < x_j - D_j \\ r_j [x_j - y_j^*] + \mathbb{E}\big\{v_{j+1}(y_j^*, D_{j+1})\big\} & \text{if } x_j - D_j \leq y_j^* \leq x_j \\ \mathbb{E}\big\{v_{j+1}(x_j, D_{j+1})\big\} & \text{if } x_j < y_j^*. \end{cases}$$

Therefore, we can compute a stochastic subgradient of $\mathbb{E}\big\{v_j(\cdot, D_j)\big\}$ at x_j through the recursion

$$\dot{v}_j(x_j, D_j) = \begin{cases} \mathbb{E}\big\{\dot{v}_{j+1}(x_j - D_j, D_{j+1})\big\} & \text{if } y_j^* < x_j - D_j \\ r_j & \text{if } x_j - D_j \leq y_j^* \leq x_j \\ \mathbb{E}\big\{\dot{v}_{j+1}(x_j, D_{j+1})\big\} & \text{if } x_j < y_j^*. \end{cases}$$

To construct tractable approximations to the stochastic subgradients of $\{f_j(\cdot) : j = 1, \ldots, n\}$, we *mimic* the computation above by using the estimates of the optimal protection levels. In particular, letting $\{y_j^k : j = 1, \ldots, n\}$ be the estimates of the optimal protection levels at iteration k and using $\mathscr{O}(\cdot)$ to denote the operator that rounds a scalar to a nearest integer by breaking ties arbitrarily, we recursively define

$$\rho_j^k(x_j, D_j, D_{j+1}, \ldots, D_n) = \begin{cases} \rho_{j+1}^k(x_j - D_j, D_{j+1}, \ldots, D_n) & \text{if } \mathscr{O}(y_j^k) < x_j - D_j \\ r_j & \text{if } x_j - D_j \leq \mathscr{O}(y_j^k) \leq x_j \\ \rho_{j+1}^k(x_j, D_{j+1}, \ldots, D_n) & \text{if } x_j < \mathscr{O}(y_j^k). \end{cases}$$

We use $\rho_j^k(x_j, D_j, D_{j+1}, \ldots, D_n)$ to approximate $\dot{v}_j(x_j, D_j)$. Specifically, at iteration k, we replace $\dot{v}_{j+1}(y_j, D_{j+1})$ in (3) with $\rho_{j+1}^k(y_j, D_{j+1}, \ldots, D_n)$ and use

$$s_j^k(y_j, D_{j+1}, \ldots, D_n) = -r_j + \rho_{j+1}^k(y_j, D_{j+1}, \ldots, D_n)$$

to approximate a stochastic subgradient of $f_j(\cdot)$ at y_j. Therefore, we propose the following algorithm to compute the optimal protection levels.

Algorithm 1

Step 1. Initialize the estimates of the optimal protection levels $\{y_j^1 : j = 1,\ldots,n\}$ such that $c \geq y_1^1 \geq y_2^1 \geq \ldots \geq y_n^1 \geq 0$. Initialize the iteration counter by setting $k = 1$.

Step 2. Letting $\{D_j^k : j = 1,\ldots,n\}$ be the demand random variables at iteration k, set

$$y_j^{k+1} = \max\Big\{ \min\big\{ \big[y_j^k + \alpha_j^k s_j^k(y_j^k, D_{j+1}^k, \ldots, D_n^k) \big]^+, c \big\}, \mathscr{O}(y_{j+1}^{k+1}) \Big\}$$

for all $j = 1,\ldots,n$.

Step 3. Increase k by 1 and go to Step 2.

Throughout, we let $\mathscr{F}^k$ be the standard filtration generated by the random variables in an algorithm. We assume that the distribution of $\{D_j^k : j = 1,\ldots,n\}$ conditional on $\mathscr{F}^k$ is the same as the distribution of $\{D_j : j = 1,\ldots,n\}$. We have the following convergence result for Algorithm 1.

Proposition 1. Assume that the sequence $\{y_j^k : j = 1,\ldots,n\}_k$ is generated by Algorithm 1. If the sequence of step size parameters $\{\alpha_j^k : j = 1,\ldots,n\}_k$ is positive and satisfies $\sum_{k=1}^{\infty} \alpha_j^k = \infty$ and $\sum_{k=1}^{\infty} [\alpha_j^k]^2 < \infty$ for all $j = 1,\ldots,n$, then we have $\lim_{k\to\infty} |y_j^k - y_j^*| = 0$ w.p.1 for all $j = 1,\ldots,n$ for some set of optimal protection levels $\{y_j^* : j = 1,\ldots,n\}$.

It is also possible to show that only knowing whether the demand strictly exceeds the number of seats that we make available for sale to each fare class is adequate to compute $\rho_j^k(x_j, D_j^k, D_{j+1}^k, \ldots, D_n^k)$. Therefore, Algorithm 1 is applicable when the demand information is censored. We also note that the way we update our estimates of the optimal protection levels ensures that we have $c \geq \mathscr{O}(y_1^k) \geq \mathscr{O}(y_2^k) \geq \ldots \geq \mathscr{O}(y_n^k) \geq 0$ for all $k = 1,2,\ldots$. Therefore, the estimates of the optimal protection levels at each iteration are nested. Finally, there is some work by van Ryzin and McGill (2000) and Huh and Rusmevichientong (2006) to compute the optimal protection levels through stochastic approximation methods, but this work assumes that the demand distributions are continuous.

2.3 Computational Results

Figure 1 compares the performance of Algorithm 1 with that of the method proposed by van Ryzin and McGill (2000) on a test problem with 12 fare classes. The two data series plot the total expected revenues corresponding to the protection levels obtained by the two methods as a function of the iteration counter. We normalize the total expected revenues by dividing by the optimal total expected revenue.

Figure 1: Performance of Algorithm 1.

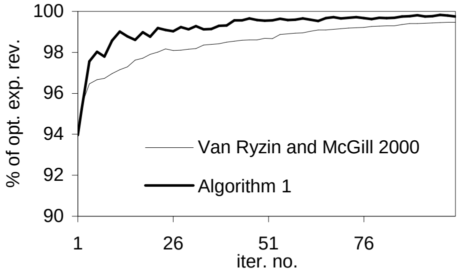

3 A Stochastic Approximation Method for the Multiple Leg Problem

We have a set of flight legs that can be used to satisfy the itinerary requests that arrive randomly over time. Whenever an itinerary request arrives, we have to decide whether to accept or reject it. An accepted itinerary request generates a revenue and consumes the capacities on the relevant flight legs. A rejected itinerary request simply leaves the system. We are interested in finding a bid price based control policy that performs well.

The idea of using bid prices as a control policy date back to Simpson (1989) and Williamson (1992). In these papers, the fundamental idea is to use the dual multipliers from a deterministic linear program as the bid prices. Talluri and van Ryzin (1998) show that using bid price policies is asymptotically optimal as the capacities on the flight legs and the expected numbers of itinerary requests increase linearly with the same rate. Recently, Bertsimas and Popescu (2003), Adelman (2006) and Topaloglu and Kunnumkal (2006) propose new mathematical programming based methods to compute bid prices. Our goal is to use stochastic approximation.

3.1 Decision and Revenue Functions

We use $\mathscr{L}$ to denote the set of flight legs and $\mathscr{J}$ to denote the set of possible itineraries. If we accept a request for itinerary j, then we generate a revenue of $\tilde{r}_j$ and consume $\tilde{a}_{ij}$ units of capacity on flight leg i. If flight leg i is not in itinerary j, then we have $\tilde{a}_{ij} = 0$. The itinerary requests arrive sequentially and we index them by $t \in \{1,2,\ldots\}$. We use x_{it} to denote the remaining capacity on flight leg i just before making the decision for itinerary request t. Therefore, the initial capacity on flight leg i is x_{i1} and $x_1 = \{x_{i1} : i \in \mathscr{L}\}$ is a part of the problem data.

Assuming that the total number of itinerary requests is bounded by a finite integer τ, we characterize the arrivals of the itinerary requests by the stochastic process $\tilde{\omega} = \{J_t : t = 1,\ldots,\tau\}$. Itinerary request t is for itinerary J_t and the value of the random variable J_t becomes known just before making the decision for itinerary request t. If we let $r_t = \tilde{r}_{J_t}$, $a_{it} = \tilde{a}_{iJ_t}$ and $a_t = \{a_{it} : i \in \mathscr{L}\}$, then we can alternatively characterize the arrivals of the itinerary requests by the stochastic process $\omega = \{(r_t, a_t) : t = 1, \ldots, \tau\}$.

The policy characterized by bid prices $\lambda = \{\lambda_i : i \in \mathscr{L}\}$ accepts an itinerary request if and only if there is enough capacity and the revenue from the itinerary request exceeds the sum of the bid prices associated with the flight legs in the requested itinerary. Therefore, as a function of the remaining leg capacities and itinerary requests, the decision function of this policy can be written as

$$\tilde{u}_t(x_t, \omega, \lambda) = \mathbf{1}(x_t \geq a_t, \, r_t \geq \textstyle\sum_{i \in \mathscr{L}} a_{it} \, \lambda_i),$$

where $\mathbf{1}(\cdot)$ is the indicator function and $x_t = \{x_{it} : i \in \mathscr{L}\}$ are the remaining leg capacities just before making the decision for itinerary request t. If the policy accepts itinerary request t, then we have $\tilde{u}_t(x_t, \omega, \lambda) = 1$. However, due to its discrete nature, the decision function above is not appropriate for developing a stochastic approximation method. Instead, we work with a smoothed version of the problem.

The fundamental idea behind the smoothed version of the problem is to assume that the leg capacities are continuous and we can accept a fraction of an itinerary request. Specifically, we consider a policy that accepts a fraction of an itinerary request depending on how much the revenue from the itinerary request exceeds the sum of the bid prices associated with the flight legs in the requested itinerary. For this purpose, we let $\theta(\cdot)$ be an increasing and differentiable function with Lipschitz derivative that satisfies $\lim_{p \to \infty} \theta(p) = 1$ and $\lim_{p \to -\infty} \theta(p) = 0$. The policy characterized by bid prices λ accepts $\theta(r_t - \sum_{i \in \mathscr{L}} a_{it} \, \lambda_i)$ units of itinerary request t as long as there is enough capacity. Furthermore, to ensure that the decision function is differentiable w.p.1, we let $\alpha = \{\alpha_{it} : i \in \mathscr{L}, \; t = 1, \ldots, \tau\}$ be uniformly distributed random variables over a small interval $[0, \varepsilon]$ and perturb the leg capacities by $\alpha_t = \{\alpha_{it} : i \in \mathscr{L}\}$ just before making the decision for itinerary request t. Therefore, we use the decision function

$$\begin{aligned} &u_t(x_t, \omega, \alpha, \lambda) \\ &\qquad = \min \left\{ \min_{i \in \mathscr{L}_t^+} \{[x_{it} + \alpha_{it}]/a_{it}\}, \, \theta(r_t - \textstyle\sum_{i \in \mathscr{L}} a_{it} \, \lambda_i) \right\}. \end{aligned}$$

It is easy to see that we have $u_t(x_t, \omega, \alpha, \lambda) \, a_{it} \leq x_{it}$ for all $i \in \mathscr{L}$ and the decision function above does not violate the leg capacities. In this case, the cumulative revenue function of the policy characterized by bid prices λ can be written recursively as

$$\begin{aligned} R_t(x_t, \omega, \alpha, \lambda) &= r_t \, u_t(x_t, \omega, \alpha, \lambda) \\ &\quad + R_{t+1}(x_t + \alpha_t - u_t(x_t, \omega, \alpha, \lambda) \, a_t, \omega, \alpha, \lambda). \end{aligned} \tag{4}$$

Since the random variable $R_1(x_1, \omega, \alpha, \lambda)$ gives the total revenue generated from all itinerary requests, we can compute a good set of bid prices by solving the problem

$$\max_{\lambda} \; \mathbb{E}\big\{R_1(x_1, \omega, \alpha, \lambda)\big\}. \tag{5}$$

3.2 Stochastic Approximation Method

It is possible to check by backward induction on (4) that $R_1(x_1, \omega, \alpha, \lambda)$ is differentiable with respect to λ w.p.1 and this derivative can easily be computed by a recursive computation. In this case, we can use the following algorithm to search for a good set of protection levels.

Algorithm 2
Step 1. Initialize the bid prices $\lambda^1 = \{\lambda_i^1 : i \in \mathscr{L}\}$ arbitrarily and initialize the iteration counter by letting $k = 1$.
Step 2. Letting ω^k be the itinerary requests and α^k be the perturbation random variables at iteration k, compute the derivative of $R_1(x_1, \omega^k, \alpha^k, \lambda^k)$ with respect to λ_i for all $i \in \mathscr{L}$. We denote this derivative by $\partial_i R_1(x_1, \omega^k, \alpha^k, \lambda^k)$.
Step 3. Letting σ^k be a step size parameter, compute the bid prices $\lambda^{k+1} = \{\lambda_i^{k+1} : i \in \mathscr{L}\}$ at the next iteration as $\lambda_i^{k+1} = \lambda_i^k + \sigma^k \, \partial_i R_1(x_1, \omega^k, \alpha^k, \lambda^k)$ for all $i \in \mathscr{L}$.
Step 4. Increase k by 1 and go to Step 2.

Assuming that the distribution of (ω^k, α^k) conditional on $\mathscr{F}^k$ is the same as the distribution of (ω, α), we have the following convergence result for Algorithm 2.

Proposition 2. Assume that the sequence of step size parameters $\{\sigma^k\}_k$ is positive and satisfies $\sum_{k=1}^{\infty} \sigma^k = \infty$ and $\sum_{k=1}^{\infty} [\sigma^k]^2 < \infty$. If the sequence of bid prices $\{\lambda^k\}_k$ is generated by Algorithm 2, then we have $\lim_{k \to \infty} \mathbb{E}\big\{\partial_i R_1(x_1, \omega, \alpha, \lambda^k)\big\} = 0$ w.p.1 for all $i \in \mathscr{L}$.

3.3 Computational Results

For four test problems, Table 1 compares the total expected revenues obtained by the bid prices computed through Algorithm 2 and the deterministic linear program proposed by Simpson (1989). The first column in this table shows the number of flight legs and the expected number of itinerary requests for each test problem. The results indicate that the bid prices computed through Algorithm 2 can perform significantly better.

Table 1: Performance of Algorithm 2.

(# legs, # itins.)	Alg. 2	Determ. LP
(12, 240)	22711	21170
(12, 240)	35603	33211
(24, 480)	36290	34008
(24, 480)	63387	55894

4 An Approximate Dynamic Programming Method for the Multiple Leg Problem

We are interested in the same multiple leg problem described in the previous section and we continue using the same notation, but for notational clarity, we use r_j and a_{ij} instead of $\tilde{r}_j$ and $\tilde{a}_{ij}$. We now assume that the different itinerary requests are independent and we let p_{jt} be the probability that itinerary request t is for itinerary j. In this case, letting e_i be the $|\mathscr{L}|$-dimensional unit vector with a 1 in the element corresponding to $i \in \mathscr{L}$, the optimal policy can be found by computing the value functions $\{V_t(\cdot) : t = 1, \ldots, \tau\}$ through the optimality equation

$$
\begin{aligned}
V_t(x_t) = \max \quad & \sum_{j \in \mathscr{J}} p_{jt} \Big[f_j \, u_{jt} + V_{t+1}(x_t - u_{jt} \textstyle\sum_{i \in \mathscr{L}} a_{ij} \, e_i) \Big] \\
\text{subject to} \quad & a_{ij} \, u_{jt} \le x_{it} \qquad \text{for all } i \in \mathscr{L}, \ j \in \mathscr{J} \\
& u_{jt} \in \{0,1\} \qquad \text{for all } j \in \mathscr{J},
\end{aligned}
$$

where u_{jt} takes value 1 if we accept itinerary request t whenever this itinerary request is for itinerary j.

4.1 Leg Based Relaxation

The multiple leg problem is difficult because if we accept an itinerary request, then we have to consume the capacity on every flight leg that is in the requested itinerary. In this section, we propose a method that is based on relaxing this requirement. In particular, we allow ourselves to individually accept or reject the flight legs that are in a requested itinerary. When we allow such *partially accepted* itineraries, the problem decomposes by the flight legs.

We begin by introducing some new notation. We augment $\mathscr{L}$ by a fictitious flight leg ψ with infinite capacity. We extend the decisions at time period t as $y_t = \{y_{ijt} : i \in \mathscr{L} \cup \{\psi\}, \ j \in \mathscr{J}\}$, where y_{ijt} takes value 1 if we accept flight leg i whenever itinerary request t is for itinerary j. In this case, it is easy to see that the optimality equation

$$
\begin{aligned}
V_t(x_t) = \max \quad & \sum_{j \in \mathscr{J}} p_{jt} \Big\{ f_j \, y_{\psi jt} + V_{t+1}(x_t - \sum_{i \in \mathscr{L}} y_{ijt} \, a_{ij} \, e_i) \Big\} \\
\text{subject to} \quad & a_{ij} \, y_{ijt} \le x_{it} \qquad \text{for all } i \in \mathscr{L}, \ j \in \mathscr{J} \\
& y_{ijt} - y_{\psi jt} = 0 \qquad \text{for all } i \in \mathscr{L}, \ j \in \mathscr{J} \\
& y_{ijt} \in \{0,1\} \qquad \text{for all } i \in \mathscr{L} \cup \{\psi\}, \ j \in \mathscr{J}
\end{aligned}
$$

is equivalent to the optimality equation given in Section 4. In the feasible solution set of the problem above, only the second set of constraints link the different flight legs. This suggests associating the dual multipliers $\alpha = \{\alpha_{ijt} : i \in \mathscr{L}, \ j \in \mathscr{J}, \ t \in \mathscr{T}\}$ with these constraints and solving the dynamic program

$$
\begin{aligned}
\tilde{V}_t(x_t \,|\, \alpha) = \max \quad & \sum_{j \in \mathscr{J}} p_{jt} \Big\{ \big[f_j - \sum_{i \in \mathscr{L}} \alpha_{ijt} \big] \, y_{\psi jt} + \sum_{i \in \mathscr{L}} \alpha_{ijt} \, y_{ijt} \\
& \qquad + \tilde{V}_{t+1}(x_t - \sum_{i \in \mathscr{L}} y_{ijt} \, a_{ij} \, e_i \,|\, \alpha) \Big\} \\
\text{subject to} \quad & a_{ij} \, y_{ijt} \le x_{it} \qquad \text{for all } i \in \mathscr{L}, \ j \in \mathscr{J} \\
& y_{ijt} \in \{0,1\} \qquad \text{for all } i \in \mathscr{L} \cup \{\psi\}, \ j \in \mathscr{J}.
\end{aligned}
$$

Letting $y_{it} = \{y_{ijt} : j \in \mathscr{J}\}$, we define the set $\mathscr{Y}_{it}(x_{it}) = \{y_{it} \in \{0,1\}^{|\mathscr{J}|} : a_{ij} \, y_{ijt} \le x_{it} \text{ for all } j \in \mathscr{J}\}$, in which case the two sets of constraints above can be succinctly written as $y_{it} \in \mathscr{Y}_{it}(x_{it})$ for all $i \in \mathscr{L}$ and $y_{\psi t} \in \{0,1\}^{|\mathscr{L}|}$. The following proposition shows that the last optimality equation decomposes by the flight legs.

Proposition 3. If $\{\vartheta_{it}(\cdot \,|\, \alpha) : t \in \mathscr{T}\}$ is the solution to the optimality equation

$$
\begin{aligned}
& \vartheta_{it}(x_{it} \,|\, \alpha) \\
& \quad = \max_{y_{it} \in \mathscr{Y}_{it}(x_{it})} \left\{ \sum_{j \in \mathscr{J}} p_{jt} \Big\{ \alpha_{ijt} \, y_{ijt} + \vartheta_{i,t+1}(x_{it} - a_{ij} \, y_{ijt} \,|\, \alpha) \Big\} \right\}
\end{aligned}
$$

for all $i \in \mathscr{L}$, then we have $\tilde{V}_t(x_t \,|\, \alpha) = \sum_{t'=t}^{\tau} \sum_{j \in \mathscr{J}} p_{jt'} \big[f_j - \sum_{i \in \mathscr{L}} \alpha_{ijt'} \big]^+ + \sum_{i \in \mathscr{L}} \vartheta_{it}(x_{it} \,|\, \alpha)$.

Therefore, we can efficiently compute $\tilde{V}_t(\cdot \,|\, \alpha)$ by concentrating on one flight leg at a time. It is also possible to show that we have $V_t(x_t) \le \tilde{V}_t(x_t \,|\, \alpha)$. Consequently, we can solve the problem $\min_\alpha \{\tilde{V}_1(x_1 \,|\, \alpha)\}$ to obtain the tightest possible upper bound on $V_1(x_1)$, which is the maximum total expected revenue generated from all itinerary requests.

4.2 Computational Results

Letting α^* be the optimal solution to $\min_\alpha \{\tilde{V}_1(x_1 \,|\, \alpha)\}$, we propose using $\tilde{V}_t(\cdot \,|\, \alpha^*)$ as an approximation to $V_t(\cdot)$. Table 2 compares the total expected revenues obtained by using $\tilde{V}_t(\cdot \,|\, \alpha^*)$ as an approximation to $V_t(\cdot)$ with those obtained by using the bid prices computed by the deterministic linear program. The results indicate that the leg based relaxation can perform significantly better.

Table 2: Performance of leg based relaxation.

(# legs, # itins.)	Alg. 2	Determ. LP
(8, 200)	18433	17873
(8, 200)	16019	15345
(16, 200)	16996	16264
(16, 200)	14753	13974

REFERENCES

Adelman, D. 2006. Dynamic bid-prices in revenue management. *Operations Research* (to appear).

Bertsekas, D. P., and J. N. Tsitsiklis. 1996. *Neuro-dynamic programming*. Belmont, MA: Athena Scientific.

Bertsimas, D., and I. Popescu. 2003. Revenue management in a dynamic network environment. *Transportation Science* 37:257–277.

Brumelle, S. L., and J. I. McGill. 1993. Airline seat allocation with multiple nested fare classes. *Operations Research* 41:127–137.

Ermoliev, Y. 1988. Stochastic quasigradient methods. In Numerical Techniques for Stochastic Optimization, ed. Y. Ermoliev and R. Wets. Berlin: Springer-Verlag.

Huh, W. T., and P. Rusmevichientong. 2006. Adaptive capacity allocation with censored demand data: Application of concave umbrella functions. Technical report, Cornell University, School of Operations Research and Industrial Engineering.

Kushner, H. J., and D. S. Clark. 1978. *Stochastic approximation methods for constrained and unconstrained systems*. Berlin: Springer-Verlang.

Simpson, R. W. 1989. Using network flow techniques to find shadow prices for market and seat inventory control. Technical report, MIT Flight Transportation Laboratory Memorandum M89-1, Cambridge, MA.

Talluri, K., and G. van Ryzin. 1998. An analysis of bid-price controls for network revenue management. *Management Science* 44 (11): 1577–1593.

Talluri, K. T., and G. J. van Ryzin. 2004. *The theory and practice of revenue management*. Kluver Academic Publishers.

Topaloglu, H., and S. Kunnumkal. 2006. Computing time-dependent bid-prices in network revenue management problems. Technical report, Cornell University, School of Operations Research and Industrial Engineering.

van Ryzin, G., and J. McGill. 2000. Revenue management without forecasting or optimization: An adaptive algorithm for determining airline seat protection levels. *Management Science* 46 (6): 760–775.

Williamson, E. L. 1992. *Airline network seat control*. Ph. D. thesis, Massachusetts Institute of Technology, Cambridge, MA.