

A Randomized Linear Programming Method for Computing Network Bid Prices

KALYAN TALLURI

Universitat Pompeu Fabra, Barcelona, Spain

GARRETT VAN RYZIN

Columbia University, New York, New York

We analyze a randomized version of the deterministic linear programming (DLP) method for computing network bid prices. The method consists of simulating a sequence of realizations of itinerary demand and solving deterministic linear programs to allocate capacity to itineraries for each realization. The dual prices from this sequence are then averaged to form a bid price approximation. This randomized linear programming (RLP) method is only slightly more complicated to implement than the DLP method. We show that the RLP method can be viewed as a procedure for estimating the gradient of the expected perfect information (PI) network revenue. That is, the expected revenue obtained with full information on future demand realizations. The expected PI revenue can, in turn, be viewed as an approximation to the optimal value function. We establish conditions under which the RLP procedure provides an unbiased estimator of the gradient of the expected PI revenue. Computational tests are performed to evaluate the revenue performance of the RLP method compared to the DLP.

A central problem in network revenue management is determining optimal decision rules for sequentially accepting or denying itinerary requests. Bid-price controls are one such class of decision rules. In a bid price control, threshold values (called bid prices) are set for each leg of a network, and an itinerary (path on the network) is accepted only if its fare exceeds the sum of the bid prices along the path. SIMPSON (1989) and WILLIAMSON (1992) first studied this method and proposed approximations to generate bid prices based on dual prices of various mathematical programming formulations of the problem.

1. INTRODUCTION AND OVERVIEW

IN TALLURI AND VAN RYZIN (1996), we analyzed the structure of an optimal policy for a dynamic stochastic model of network revenue management. The same model is used in this paper. Time is discrete. We consider an m -leg network, and let $x = (x^1, \dots, x^m)$ denote the (integer) vector of remaining leg capacities and k denote the number of time periods

remaining. As a convention, we use superscripts to index components of a vector and use subscripts to denote time or elements of a sequence. We use x^T to denote the transpose of a vector x .

As in Talluri and van Ryzin (1996), we model demand as a sequence of requests for itineraries (paths) over time. An itinerary consists of a collection of legs together with an associated fare. To model multiple fare classes, we define several itineraries (one for each fare class), each having an identical set of legs but different fares. As a simplification, we assume itineraries that are rejected are lost to the network. In particular, we do not model diversion among itineraries.

Define $A = [a_{ij}]$ where $a_{ij} = 1$ if itinerary j uses leg i and $a_{ij} = 0$ otherwise; the j th column of A , A^j , is the incidence vector for itinerary j . Abusing this notation somewhat, we shall also use $i \in A^j$ to indicate the legs i that are used by itinerary j . We assume A^j has at least one nonzero component, i.e., all itineraries use at least one leg. Each itinerary j contributes revenue r^j . Define $r = (r^1, \dots, r^n)$.

In this paper, we will not need to specify the

arrival process in detail; the only demand information we need is the distribution of total demand to come at each point in time. However, as a canonical example of an arrival process, one can consider the case where only one itinerary arrives in a period and the probability that itinerary j arrives in period k is p_k^j , where $\sum_{j=1}^n p_k^j \leq 1$ for all k . Let Y^j denote the total number of requests for itinerary j in the remaining time k , and define $Y = (Y^1, \dots, Y^n)$. To avoid excessive notation, we do not explicitly index Y with the remaining time k ; however, Y will always represent the total demand to come over the remaining k units of time.

Let $J_k(x)$ denote the value function (optimal expected revenue given x and k). This function is defined formally for a general network model with Markovian arrivals in Talluri and van Ryzin (1996). However, for our purposes, an informal definition of $J_k(x)$ will suffice. In Talluri and van Ryzin (1996), we showed that it is optimal to accept a reservation for itinerary j with a revenue r^j in period k if and only if

$$r^j \geq J_{k-1}(x) - J_{k-1}(x - A^j). \quad (1)$$

This condition is intuitive: $r^j + J_{k-1}(x - A^j)$ is the expected revenue if we accept the reservation, and $J_{k-1}(x)$ is the expected revenue if we do not; therefore, if the former is greater than the latter, it makes sense to accept the reservation.

A bid price control, as defined in Talluri and van Ryzin (1996), specifies an m -vector of values (bid prices), $\mu_k(x)$, for each remaining capacity x and remaining time k , such that an itinerary j is accepted if and only if

$$r^j > \sum_{i \in A^j} \mu_k^i(x).$$

That is, itinerary j is accepted if and only if its revenue exceeds the sum of the bid prices on the legs it uses. This type of control is appealing on both intuitive and practical grounds. Intuitively, bid prices represent the marginal value of capacity (also called displacement cost) on each leg; if the current request exceeds the sum of the expected marginal values of the legs it requires, then it makes intuitive sense to accept it, because the current revenue is larger than our estimate of the next best use for the same capacity. Bid price controls are also appealing from an implementation standpoint, because they require only a small number of control values (one bid price for each leg), and the decisions to accept or reject arriving requests are trivial to determine given these values.

However, in general, the optimality condition, Eq.

1, does not lead to a bid price control. (See Talluri and van Ryzin (1996) for counter examples.) At the same time, it is easy to see how condition 1 motivates a bid price structure. Indeed, if we imagine x is a real vector and $J_{k-1}(x)$ is a differentiable function of x , then a first-order approximation to Eq. 1 yields

$$\begin{aligned} r^j &\geq J_{k-1}(x) - J_{k-1}(x - A^j) \\ &\approx \nabla_x^T J_{k-1}(x) A^j = \sum_{i \in A^j} \frac{\partial}{\partial x^i} J_{k-1}(x). \end{aligned} \quad (2)$$

If this first-order approximation is good, then using the gradient $\nabla_x J_{k-1}(x)$ as a vector of bid prices should produce close-to-optimal decisions. Although this reasoning is very informal, it lies at the heart of the concept of bid price controls. Moreover, it was proved in Talluri and van Ryzin (1996) that a bid price policy is, in fact, asymptotically optimal as the demand volume and leg capacities tend to infinity, which provides theoretical support for the use of bid price controls.

1.1 Approximation Methods

To implement a bid price control strategy, one needs a method for generating bid price values. Computing $J_k(\cdot)$ is not feasible in practice because of the enormous size of the state space. (The state space is of size $O(C^m)$, where C is the maximum leg capacity and m is the number of legs. Typical numbers for a major airline network are $C = 300$ and $m = 5000$ or more.) Therefore, the only realistic option in practice is to use approximation methods.

Such methods can generally be interpreted as consisting of two steps: 1) approximate the value function $J_k(x)$; and 2) use the associated gradient (or subgradient) information from the approximation as a vector of bid prices. That is, form an approximation $J_k^A(x)$ to the value function and use $\nabla J_k^A(x)$ (or a subgradient if $\nabla J_k^A(x)$ does not exist) for the vector of bid prices. Typically, these approximations are re-computed periodically throughout the booking horizon in response to changes in the remaining capacity or demand forecast.

Perhaps the simplest example of such an approximation is the deterministic linear programming (DLP) method, which was initially investigated by Williamson (1992). The DLP approximation to the value function is obtained by finding a tentative itinerary allocation, y , that solves

$$J_k^{\text{LP}}(x) = \max r^T y \quad (3)$$

$$Ay \leq x \quad (4)$$

$$0 \leq y \leq EY. \quad (5)$$

Note this approximation corresponds to assuming that demand to come, Y , is always equal to its mean, EY . The primal itinerary allocation, y , is not used directly. Rather, let $\mu = (\mu^1, \dots, \mu^m)$ denote the dual prices associated with constraint (4). Provided the linear program is not degenerate at the optimal solution, $\nabla J_k^{LP}(x)$ exists and is equal to μ . Therefore, the vector of dual prices μ is used for the bid prices.

Although this method is quite simple and intuitive, it has several obvious weaknesses. First, it only uses the mean demand and ignores all other distributional information on Y . This is clearly unrealistic, because it is known that, for single leg problems, the expected marginal value of seat capacity depends on the entire distribution of future demand BELOBABA (1989), BRUMELLE and MCGILL (1993), LITTLEWOOD (1972), ROBINSON (1991), and WOLLMER (1992). As a result, the DLP can produce poor approximations to the true marginal expected value. For example, one can easily show that, if the expected demand on a leg is strictly less than the capacity, then the deterministic linear program will return a zero bid price for the leg. Yet, with highly variable demand, the expected marginal value can be much higher than zero—higher than the fares of some of the lower fare classes perhaps—even though the mean demand is less than capacity. Such discrepancies can lead to poor accept/deny decisions.

However, the linear programming method is appealing precisely because it is so simple and computationally efficient. It is desirable, therefore, to have a method that is nearly as simple and efficient as the deterministic linear programming method but which also captures more distributional information about future demand.

1.2 A Randomized Linear Programming Method

One idea for incorporating more stochastic information into the linear programming method is to replace the expected value of Y in constraint (5) by the random vector Y itself. Then, the expected value of the resulting optimal solution can be used as an approximation to the value function. That is, define

$$v_k(x, Y) = \max r^T y \tag{6}$$

$$Ay \leq x \tag{7}$$

$$0 \leq y \leq Y. \tag{8}$$

The optimal value $v_k(x, Y)$ is a random variable. Let $\mu(x, Y)$ denote an optimal vector of dual prices for the set of constraints (7), and note that $\mu(x, Y)$ is also a random vector.

Next, consider the approximation to the value function,

$$J_k^{PI}(x) = E v_k(x, Y). \tag{9}$$

We call this the perfect information (PI) approximation, because it corresponds to a case in which future allocations (and revenues) are based on perfect knowledge of Y ; however, at time k , Y is not yet realized. Assuming the gradient exists, we then use $\nabla_x E v_k(x, Y)$ as our vector of bid prices.

However, this method is viable only if we can efficiently compute $\nabla_x E v_k(x, Y)$. One appealing approach is to consider interchanging differentiation and expectation. Assuming such an interchange is justified (sufficient conditions are given in Lemma 1 below), we have

$$\nabla_x E v_k(x, Y) = E \nabla_x v_k(x, Y). \tag{10}$$

This interchange, in turn, suggests a procedure for estimating $\nabla_x E v_k(x, Y)$. Simply simulate N independent samples of the demand vector, Y_1, \dots, Y_N , and solve Eqs. 6–8 for each sample. Then estimate the gradient using

$$\frac{1}{N} \sum_{i=1}^N \mu(x, Y_i). \tag{11}$$

That is, simply average the dual prices from the PI allocation over a series of randomly generated demands. We call this idea the randomized linear programming (RLP) method.

The RLP method has several appealing features. First, it is a simple modification to the DLP method, so it can be easily incorporated into production revenue management systems that use DLP. Second, it has the flexibility to model very general demand distributions, because one only needs the ability to simulate demand to apply the method (e.g., one could allow for different coefficients of variation and/or correlations among the components of Y). Finally, it incorporates distributional information on future demand.

To see why, consider the simple case where $m = 1$ (a single leg) and $A = [1]$ (one itinerary). In this case, r, x, Y , and μ are all scalars, and it is easy to see that $\mu(x, Y) = r$ if $Y > x$ and $\mu(x, Y) = 0$ if $Y < x$, so if Y has a continuous distribution, then

$$\frac{d}{dx} E v_k(x, Y) = E \mu(x, Y) = rP(Y > x),$$

which is Littlewood’s (1972) expression for the expected marginal value of capacity. That is, in this simple case, the RLP method produces the correct expected marginal value. In contrast, the DLP

method produces a marginal value of zero if $EY < x$ and a marginal value of r if $EY > x$. This example suggests why the RLP method may produce a somewhat better approximation of the marginal value of capacity than the DLP method.

It turns out that the RLP method, in fact, received some brief attention over a decade ago in a study of network control methods at American Airlines conducted by SMITH and PENN (1988). However, the idea is not widely known, and we are not aware of any extensive study of its theoretical properties or practical performance.¹ In the remainder of the paper, we provide such a study. In particular, in Section 2, we give sufficient conditions under which both Eq. 10 is justified and Eq. 11 is an unbiased (and consistent) estimator of the gradient of the expected PI revenue. These conditions also guarantee the existence of $\nabla_x Ev_k(x, Y)$. In Section 3, we examine the situation where $\nabla_x Ev_k(x, Y)$ fails to exist. In Section 4, we then perform a series of computational tests to assess the revenue performance of the RLP method relative to the DLP method. Our conclusions are presented in Section 5.

2. PROPERTIES OF THE RLP ESTIMATOR

WE NEXT ESTABLISH conditions under which the RLP estimator gives an unbiased estimate of the gradient of $Ev_k(x, Y)$. This is important because, in general, $E\mu(x, Y)$ is only a subgradient of $Ev_k(x, Y)$. If $E\mu(x, Y)$ is not a gradient, then the logic derived from Eq. 2 for using $E\mu(x, Y)$ as a vector of bid prices fails because the directional derivative (if it exists),

$$D(x; d) = \lim_{h \rightarrow 0} \frac{1}{h} [Ev_k(x + hd, Y) - Ev_k(x, Y)], \quad (12)$$

is no longer given by $E\mu^T(x, Y)d$. The following two conditions, however, are sufficient to both justify Eq. 10 and to guarantee that Eq. 11 is an unbiased estimator of the gradient:

CONDITION 1. *If $x = \sum_{j \in S} \alpha^j A^j$, then $\{A^j: j \in S\}$ must have rank m . That is, x does not lie in any subspace defined by a subset of fewer than m columns of A .*

CONDITION 2. *$P(Y^j \leq y)$ is continuous in y , for all $j = 1, \dots, n$.*

¹Smith and Penn (1988) concluded that the RLP method was too time consuming relative to the improvement it provided, and they focused their testing on the DLP method.

(Examples of failures of these conditions are provided in Section 2.2.) Our main result is summarized in Theorem 1 below. The proof of the theorem requires a sequence of lemmas, which we present next.

2.1 Proof of Unbiasedness

First, we require a lemma by GLASSERMAN (1994), which provides sufficient conditions for interchanging expectation and differentiation.

LEMMA 1. (Glasserman, 1994). *Let $X(\theta)$ be a random function satisfying the Lipschitz condition: There exists a $K \geq 0$ such that*

$$|X(\theta + h) - X(\theta)| \leq Kh, \quad \text{for all } h \geq 0.$$

Suppose that $X(\theta)$ is differentiable at θ (a.s.). Then $(d/d\theta)EX(\theta)$ exists and

$$\frac{d}{d\theta} EX(\theta) = E \frac{d}{d\theta} X(\theta).$$

The first condition of the lemma is easy to establish for the RLP method. Indeed, it is easy to see that if $\bar{r} = \max\{r^j: j = 1, \dots, n\}$ and e^i is the i th unit vector, then, from Eqs. 6–8 and the fact that elements of A are 0 or 1, $|v_k(x + he^i, Y) - v_k(x, Y)| \leq h\bar{r}$ for all x and i . That is, increasing x^i by h allows at most an increase of h in $\sum_{j=1}^n y^j$, which contributes at most $\bar{r}h$ to the objective function value. Therefore, we have

LEMMA 2. *The function $v_k(x, Y)$ defined in Eqs. 6–8 satisfies the Lipschitz condition in Lemma 1.*

From Lemma 1, it then follows that if $\nabla_x v_k(x, Y)$ exists with probability one, Eq. 10 is justified. Note that $\nabla_x v_k(x, Y)$ exists if and only if the dual price $\mu(Y)$ is unique. Therefore, we must establish conditions under which $\mu(Y)$ is unique with probability one.

To do so, consider the Lagrangian dual of the linear program, Eqs. 6–8,

$$\begin{aligned} L(x, \mu) &= \max_{0 \leq y \leq Y} \sum_{j=1}^n y^j (r^j - \mu^T A^j) + \mu^T x \\ &= \sum_{j=1}^n Y^j (r^j - \mu^T A^j)^+ + \mu^T x, \end{aligned} \quad (13)$$

where $(x)^+ = \max\{0, x\}$. Note that $L(x, \mu)$ is convex in μ . Let μ_* denote a solution to the problem

$$\min_{\mu \geq 0} L(x, \mu). \quad (14)$$

By strong duality, $v_k(x, Y) = L(x, \mu_*)$ and μ_* is always a subgradient of $v_k(x, Y)$. If μ_* is unique, then it is also a gradient, and $\nabla_x v_k(x, Y)$ exists. [See Chapter 6 of Bazaraa, Sherali, and Shetty (1993).] Therefore, we must show that μ_* is unique with probability one.

To this end, let $d \neq 0$ denote a feasible direction from μ_* , i.e., there exists an $\epsilon > 0$ such that

$$\mu_* + hd \geq 0, \quad h \in [0, \epsilon]. \quad (15)$$

Define the sets

$$E_0 = \{j: r^j - \mu_*^T A^j > 0\}, \quad (16)$$

$$E_1 = \{j: r^j - \mu_*^T A^j = 0, r^j - \mu_*^T A^j - hd^T A^j > 0\}, \quad (17)$$

$$E_2 = \{j: r^j - \mu_*^T A^j = 0, r^j - \mu_*^T A^j - hd^T A^j \leq 0\}, \quad (18)$$

and note from Eq. 13 that

$$L(x, \mu_* + hd) - L(x, \mu_*) = hd^T \left(x - \sum_{j \in E_0 \cup E_1} Y^j A^j \right). \quad (19)$$

The solution μ_* is unique if the right hand side above is strictly positive for all feasible directions d .

We need the following lemma.

LEMMA 3. *Suppose that C1 holds, $d^T A^j = 0, \forall j \in E_0 \cup E_1$, and $d^T x = 0$. Then $d = 0$.*

Proof. By Eq. 19 and the optimality of μ_* , it must be true that the system,

$$\begin{aligned} \lambda^T A^j &= 0, \quad j \in E_0 \cup E_1 \\ \lambda^T x &< 0, \end{aligned}$$

is unsolvable, else λ is a strictly improving direction. Thus, by Farka's lemma, the system

$$\sum_{j \in E_0 \cup E_1} \alpha^j A^j = x \quad (20)$$

is solvable. But, by condition C1, $\{A^j: j \in E_0 \cup E_1\}$ has rank m , so $d^T A^j = 0, \forall j \in E_0 \cup E_1$ implies $d = 0$. \square

The almost sure existence of the gradient is established in the next lemma.

LEMMA 4. *Suppose C1 and C2 hold. Then $\nabla_x v_k(x, Y)$ exists with probability one.*

Proof. As mentioned above, it suffices to show that μ_* is a unique optimal solution to Eq. 14 with probability one. To do so, we must show that, with

probability one, there are no directions $d \neq 0$ for which $L(x, \mu + hd) = L(x, \mu)$. The proof is by contradiction.

Indeed, if $d \neq 0$ and $L(x, \mu + hd) = L(x, \mu)$, then, from Eq. 19, we must have

$$d^T \left(x - \sum_{j \in E_0 \cup E_1} Y^j A^j \right) = 0.$$

By Condition 2, it is clear that $P(x = \sum_{j \in E_0} \cup E_1 Y^j A^j) = 0$. Hence, the only way the above equality can hold with positive probability is if $d^T x = 0$ and $d^T A^j = 0, j \in E_0 \cup E_1$. But then, Lemma 3 implies $d = 0$, a contradiction. \square

THEOREM 1. *Suppose conditions C1 and C2 hold, then $\nabla_x E v_k(x, Y)$ exists and*

$$\nabla_x E v_k(x, Y) = E \nabla_x v_k(x, Y) = E \mu(x, Y).$$

Proof. This follows directly from Lemmas 1, 2, and 4. \square

Theorem 1 shows that, under C1 and C2, the RLP estimator, Eq. 11, is indeed an unbiased estimator of the gradient $\nabla_x E v_k(x, Y)$; by the strong law of large numbers, the estimator is also consistent.

2.2 Examples of Nondifferentiability

To gain some insight into Conditions 1 and 2, it is useful to consider simple examples in which the linear program, Eqs. 6–8 is degenerate with positive probability, and thus $\nabla_x E v_k(x, Y)$ does not exist.

First, consider the following example which violates Condition 1.

$$r = (100, 1, 1) \quad x = (2, 2)$$

$$A = \begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}.$$

Note $x = 2A^1$, so Condition 1 is violated as claimed. Assume Condition 2 holds, however, and that $P(Y^1 > 2) = 1/2$. Then, for realizations with $Y^1 > 2$, it is clear that the optimal solution is $y = (2, 0, 0)$, a degenerate optimal solution. Therefore, the linear program is degenerate with probability at least $1/2$.

Next, suppose

$$r = (100, 1000, 1) \quad x = (2, 3)$$

$$A = \begin{bmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix},$$

so Condition 1 is satisfied, but suppose $P(Y^2 = 1) = 1/2$, so Condition 2 is violated. Also, assume $P(Y^2 > 2) = 1/2$. Then, when $Y^2 = 1$ and $Y^2 > 2$, the linear program has optimal solution $x = (1, 1, 0)$, which is

degenerate. Therefore, the linear program is degenerate with probability at least $\frac{1}{4}$.

3. SOME COMMENTS ON THE NONDIFFERENTIABLE CASE

LET $\mathcal{D} \subset R^n$ DENOTE the set of demand vectors Y for which the linear program, Eqs. 6–8, is degenerate. Let $\mathcal{D}^c = R^n - \mathcal{D}$, i.e., \mathcal{D}^c is the complement of \mathcal{D} . The nondifferentiable case occurs when $P(Y \in \mathcal{D}) > 0$. In this section, we discuss some issues related to this case.

3.1 Convergence to a Subgradient Estimator

If the linear program, Eqs. 6–8, is degenerate for some Y , then the dual solution $\mu(x, Y)$ will depend on the solution algorithm (e.g., interior point methods may return a different dual solution than simplex-based methods). Thus, we consider $\mu(x, Y)$ to be defined as the dual solution returned by a given linear programming algorithm.

Note that, for any realization of demand Y , $\mu(x, Y)$ is a subgradient of $v_k(x, Y)$, i.e.,

$$v_k(z, Y) \geq v_k(x, Y) + \mu^T(x, Y)(z - x)$$

for all z . Because all terms above are uniformly bounded in Y , taking expectations on both sides implies

$$J_k^{\text{PI}}(z) \geq J_k^{\text{PI}}(x) + E\mu^T(x, Y)(z - y),$$

and hence $E\mu(x, Y)$ is a subgradient of $J_k^{\text{PI}}(x)$.

However, the existence of $\lim_{N \rightarrow \infty} (1/N) \sum_{i=1}^N \mu(x, Y_i)$ will depend on how Eqs. 6–8 is solved at each step i . For example, it is possible that the linear programming algorithm is initialized with a previous solution (warm started), and hence $\mu(x, Y_i)$ becomes a function of the history of the sequence of demand vectors $\{Y_1, \dots, Y_{i-1}\}$. [Our experience in numerical testing is that such warm starts significantly speed up the computation of $\mu(x, Y_i)$.] Under such conditions, it is possible that the estimator, Eq. 11 does not converge. However, the following proposition is clearly true.

PROPOSITION 1. *Let A denote the linear programming algorithm used to solve Eqs. 6–8. If algorithm A is initialized with the same data on each run (e.g., it is not warm started with data from prior solutions), then $\lim_{N \rightarrow \infty} (1/N) \sum_{i=1}^N \mu(x, Y_i)$ always exists and is a subgradient of $E v_k(x, Y)$.*

3.2 Directional Derivatives

As mentioned, when $E\mu(x, Y)$ is only a subgradient, then the logic derived from Eq. 2 for using $E\mu(x, Y)$ as a vector of bid prices fails because the

directional derivative (if it exists) is no longer given by $E\mu^T(x, Y)d$. We next show that the directional derivatives $D(x; d)$ (see Eq. 12) always exist. The proof leads to a natural algorithmic modification to the estimator, Eq. 11, to make it an unbiased estimator of directional derivatives.

Define $u(x, Y; d)$ as the optimal value of the linear program

$$u(x, Y; d) = \min d^T \mu \quad (22)$$

$$\mu^T x = v_k(x, Y)$$

$$\mu^T A \leq r$$

$$\mu \geq 0.$$

Then we have

PROPOSITION 2. *The directional derivative, $D(x; d)$, exists and*

$$D(x; d) = E u(x, Y; d).$$

Proof. Let $\partial v_k(x, Y)$ denote the set of subgradients of $v_k(x, Y)$ and note that $\partial v_k(x, Y) = \{\mu: \mu^T x = v_k(x, Y), \mu^T A \leq r, \mu \geq 0\}$. Observe that $\partial v_k(x, Y)$ is uniformly bounded for all Y , provided there are no zero columns of A . It is well known (Bazaraa, Sherali, and Shetty (1993), page 218.) that the directional derivative, $D(x, Y; d)$, satisfies

$$\begin{aligned} D(x, Y; d) &= \lim_{h \rightarrow 0} \frac{1}{h} [v_k(x + hd, Y) - v_k(x, Y)] \\ &= \inf\{d^T \mu: \mu \in \partial v_k(x, Y)\}. \end{aligned}$$

Note the right hand side above is simply $u(x, Y; d)$. Because $\partial v_k(x, Y)$ is uniformly bounded for all Y , it follows that the $D(x, Y; d)$ exists for all Y and is also uniformly bounded. Applying Lemma 1 establishes the result. \square

Proposition 2 shows that, by solving the auxiliary linear program, Eq. 22 for each sample, we can obtain the unbiased estimator of the directional derivative,

$$\frac{1}{N} \sum_{i=1}^N u(x, Y_i; d).$$

The disadvantage of this approach, of course, is that this estimator is for a single direction d . To make effective use of such information in network revenue management, we need to evaluate the derivative for every itinerary j (i.e., every direction A^j) by solving a separate sequence of linear programs. For a large airline network, the number of itineraries, n , can be

on the order of 100,000, making such an approach somewhat impractical.

4. SIMULATION EXPERIMENTS

OUR ANALYSIS THUS FAR provides a formal basis for interpreting the estimator, Eq. 11, and the RLP method in general. However, what ultimately matters in practice is how well the RLP method performs on real-world networks. In this section, we describe a series of simulation experiments comparing the revenue performance of the DLP and RLP methods. We first describe the simulation methodology and the test networks used in our tests. We then present and discuss the numerical results. Our conclusion is that the RLP method provides a small but significant improvement over DLP and, thus, warrants further consideration.

4.1 Simulation Methodology

To evaluate the performance of the RLP method, we conducted a series of simulation tests using both real-world and randomly generated networks. In our simulations, each request is for one seat (no group bookings) and we did not simulate cancellations or no-shows (no overbooking effects). Each day consists of 1000 minutes (roughly 18 hours) of booking time. Booking requests are generated over a booking horizon of approximately 20 days.

Booking requests are randomly generated for each instance using a two-step process. First, the final demand for an itinerary is generated based on a given distribution. We used both a Poisson distribution and a truncated normal distribution with a fixed coefficient of variation in our tests. Each itinerary has an associated booking curve, that specifies the fraction of total demand observed as a function of time. In the second step, this booking curve is used as a probability distribution to determine the arrival time of a request. That is, for each one of the generated requests, we generate a uniform random number between 0 and 1, and the arrival time of the request is then determined from the inverse of the booking curve function. Note that this process produces a (nonhomogeneous) Poisson process when the final distribution is Poisson; for non-Poisson cases, it produces demand that statistically conforms to both the final demand distribution and the booking curve.

The resulting stream of itinerary requests is then sorted based on their arrival times and processed sequentially by a simulated reservation system that uses a bid-price control rule to make each accept/deny decision. Bid prices are constant for each simulated day, but are recalculated at the start of each

simulated day by rerunning the bid price optimization (either DLP or RLP) using updated forecasts and capacities. This procedure mimics the overnight processing performed by airlines in practice.

Although actual arrivals were generated according to these discrete distributions, within the RLP we used continuous, truncated-normal distributions. That is, each component of Y in Eq. 8 within our RLP computations was modeled as a truncated-normal random variable with the appropriate mean and standard deviation. We did this because continuous distributions better match the assumptions of the linear programming model (e.g., continuous allocation of continuous capacity) and they also tend to produce less degeneracy.

To make the test more realistic, we also simulated a simple forecasting process. That is, the optimization did not use actual demand means and variances; rather, estimates of the means and variances were determined based on observations of past simulated demand. In this way, we tried to mimic the effect of forecast error on the control policy.

Forecasts were obtained as follows. Twenty files of arrival histories were generated according to the given distributions, corresponding to twenty past departures of each itinerary. These twenty histories were used to estimate means and variances for the first simulation run. Then, a new demand process was simulated and added to the history. These twenty-one history files were then used to estimate means and variances for the second simulation run, and so on. This forecasting methodology basically assumes stationarity and independence of demand. Given that there is no seasonality or trend in the sequence of historical files being generated, this is an accurate forecasting method. In practice, forecasting methods would need to take into account nonstationarities in the demand process.

This process was used to generate 20 simulated departure days of each network under both the RLP and DLP policy. The same sample paths of data were used for each method (coupled simulations) to reduce the variance in our performance comparisons. It appeared that 20 runs were sufficient to get a reasonably accurate estimate of the revenue differences, because increasing the number of runs showed little difference in the relative performance of the methods.

For our RLP implementation, 30 samples of the dual prices $\mu(Y)$ were averaged on each call to the optimization module. Based on some preliminary testing, it did not seem that the performance of the RLP method improved much with larger samples, though this behavior could very well be problem dependent.

4.2 Test Networks

We used three different test networks for our simulations. The first one (Network 1) is based on data from a hub-bank (a set of arriving flights connecting to a set of departing flights) of a major international airline with 102 flight legs, 1066 itineraries, 7 fare classes, and 1 compartment. The forecast of the demand means for each itinerary (called an origin-destination fare class, or ODF) were given along with capacities for each flight leg. Booking curves based on historical observations were provided for each itinerary.

From these data, we generated two different demand scenarios: one with Poisson demand for each ODF and another in which demand for each ODF has a rounded and truncated normal distribution (truncated by 0 on the left) with a coefficient of variation (CV) of 1.414. The parameters of the normal distribution are determined so that the mean and the standard deviation of its truncated version match the desired mean and standard deviation.

The second dataset (Network 2) comes from a large U.S. airline’s domestic hub-bank with 62 legs, 517 itineraries, 11 fare classes, and 1 compartment. The fare structure for this network was quite different from Network 1, exhibiting significantly lower fare dispersion. Again, both Poisson and truncated normal demand were simulated for this network.

Finally, we generated a random airline network with a hub-and-spoke topology (Network 3), consisting of 20 legs, 120 itineraries, and 8 fare classes. The demand means for each ODF were randomly generated, uniformly spread between 0.0 and 3.0. The fares for each ODF were also randomly generated from a uniform distribution in specific ranges for each fare class. The fares varied between \$30 and \$1,100. The booking curves enforced a strict low-before-high fare pattern for each O-D pair. Only the Poisson demand case was simulated for Network 3.

For each network, various load factors were simulated by scaling the mean demands by a constant. We generated load factors approximately between 60 and 85%.

For each run, an upper bound on the maximum achievable revenue was also calculated by solving the linear programming relaxation of the so-called perfect-hindsight integer program. (That is, the linear program that allocates capacity based on perfect information on the total realized demand.) This bound is likely to be loose for two reasons. First, it assumes perfect information on future demand, which is clearly overly optimistic. Second, for a hub-and-spoke network, there may be a gap between the optimal value of the integer program and its linear-




Fig. 1. Network 1 with Poisson demand: percent improvement over DLP.

programming relaxation. Nevertheless, the perfect-hindsight upper bound does give some measure of the absolute optimality gap.

4.3 Results

The results of our simulation experiments are shown in Figs. 1–4. These graphs show both the RLP and perfect-hindsight upper bound revenue, expressed as a percentage difference from the DLP revenue. The numbers in parentheses are the revenue upper bounds expressed as a percentage difference from the DLP revenue.

The RLP generates a slight improvement in revenue over the DLP policy for the two airline networks (Networks 1 and 2). The revenue improvement




Fig. 2. Network 1 with demand from a truncated normal distribution (CV = 1.414): percent improvement over DLP.

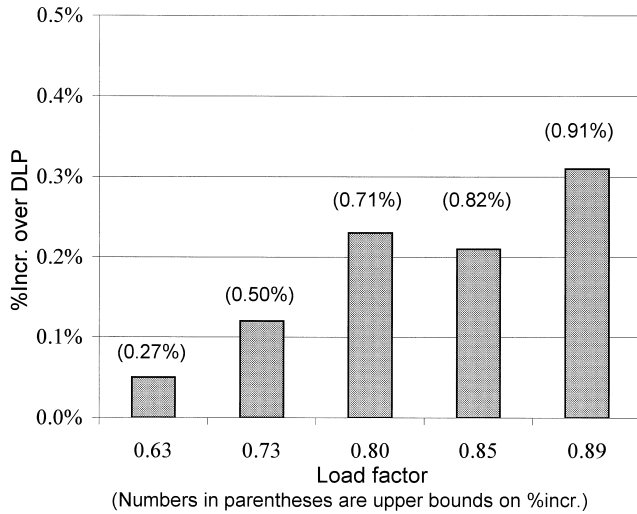


Fig. 3. Network 2 with Poisson demand: percent improvement over DLP.

ranges from 0.04 to +0.32%. Although small in absolute terms, these improvements are significant, especially at the higher load factors.

The performance of the RLP method on the randomly generated network, however, is somewhat erratic and does not uniformly dominate the DLP method. However, it seems to be at least comparable to DLP. One factor that could explain such behavior on this network is the difference in the arrival process of bookings. The real-world datasets from Networks 1 and 2 have booking curves that are closer to being uniform over time. In the random network (Network 3), bookings were generated using a strict low-before-high fare booking pattern. It is plausible that the expected perfect information revenue is a




Fig. 4. Network 2 with demand from a truncated normal distribution (CV = 1.414): percent improvement over DLP.




Fig. 5. Random Network 3 with Poisson demand: percent improvement over

better approximation of the value function in the former case. That is, assuming perfect future allocations is a particularly optimistic assumption in the low-before-high case.

Finally, we note that, on these three networks, we almost never encountered degeneracy of the linear programs of the RLP method. We attribute this to the continuous distributions used in our RLP computations combined with the irregular network capacities and itineraries on these networks. Again, however, such behavior is likely to be highly problem dependent.

5. CONCLUSIONS

THE RLP METHOD is a simple and appealing alternative to DLP for large-scale network revenue management applications. Overall, our numerical tests indicate that RLP provides a small but significant improvement over DLP, though this conclusion is preliminary and warrants further investigation.

It would be interesting to see if there are other variations of this approach that would improve revenue performance further. Also, there may be some interesting work in using variance reduction techniques to improve the efficiency of the estimator, Eq. 11. More generally, simulation-based optimization techniques may prove to be a fruitful approximation approach for future research in this area.

REFERENCES

M. S. BAZARAA, H. D. SHERALI AND C. M. SHETTY, *Nonlinear Programming: Theory and Algorithms*, John Wiley, New York, 1993.

- P. P. BELOBABA, "Application of a Probabilistic Decision Model to Airline Seat Inventory Control," *Opns. Res.* **37**, 183–197 (1989).
- S. L. BRUMELLE AND J. I. MCGILL, "Airline Seat Allocation with Multiple Nested Fare Classes," *Opns. Res.* **41**, 127–137 (1993).
- P. GLASSERMAN, "Perturbation Analysis of Production Networks," in *Stochastic Modeling and Analysis of Manufacturing Systems*, Yao, D. (ed.), Springer-Verlag, Berlin, 1994.
- K. LITTLEWOOD, "Forecasting and Control of Passengers," *12th AGIFORS Symposium Proceedings*, 95–128 (1972).
- L. W. ROBINSON, "Optimal and Approximate Control Policies for Airline Booking with Sequential Nonmonotonic Fare Classes," *Opns. Res.* **43**, 252–263 (1991).
- R. W. SIMPSON, "Using Network Flow Techniques to Find Shadow Prices for Market and Seat Inventory Control," MIT Flight Transportation Laboratory Memorandum M89-1, Cambridge, Massachusetts, 1989.
- B. C. SMITH AND C. W. PENN, "Analysis of Alternative Origin–Destination Control Strategies," in *AGIFORS Symposium* proc. 28, Airline Group of the International Federation of Operational Research Societies, 1988.
- K. T. TALLURI AND G. J. VAN RYZIN, "An Analysis of Bid-Price Controls for Network Revenue Management," *Management Sci.* **44**, 1577–1593 (1996).
- E. L. WILLIAMSON, "Airline Network Seat Control." Ph.D thesis, MIT, Cambridge, Massachusetts, 1992.
- R. D. WOLLMER, "An Airline Seat Management Model for a Single Leg Route when Lower Fare Classes Book First," *Opns. Res.* **40**, 26–37 (1992).

(Received: May 1997; revisions received: February 1998, January 1999; accepted: January 1999)