Refined Experimental Design and Regression Splines Method for Network Revenue Management

Sheela Siddappa[*,†], Dirk Günther[†,‡], Jay M. Rosenberger[§,¶], Victoria C. P. Chen[§]

June 14, 2007

COSMOS Technical Report 06-04

Sheela Siddappa *is a Researcher at General Motors in Bangalore, India. She holds a PhD in industrial engineering from The University of Texas at Arlington. Her research interests include revenue management, statistical regression, and optimization under uncertainty.*

Dirk Günther *is a Principal at Sabre Holdings in Witten, Germany. He holds a PhD in industrial engineering from the Georgia Institute of Technology. His research interests include revenue management and other applications of operations research in transportation and logistics.*

Jay Rosenberger *is an assistant professor of industrial engineering at The University of Texas at Arlington. He holds a PhD in industrial engineering from the Georgia Institute of Technology. His research interests include mathematical programming and simulation in transportation, defense, and health care. Dr. Rosenberger's graduate research on airlines won the First Place*

---

[*] General Motors Technical Centre -India 3rd Floor, Creator Building, ITPL, Whitefield Road Bangalore, Karnataka, 560066 India

[†] This research was done as part of Günther's and Siddappa's Ph.D. theses (Günther 1998; Siddappa 2006).

[†] This research was done as part of Günther's and Siddappa's Ph.D. theses (Günther 1998; Siddappa 2006).

[‡] Sabre Research Group, Bebelstrasse 15, 58453 Witten, GERMANY.

[§] Department of Industrial and Manufacturing Systems Engineering, The University of Texas at Arlington Campus Box 19017 Arlington, TX 76019 USA

[¶] Corresponding author: Tel: +1 817 272 5787; Fax: +1 817 272 3406; E-mail: jrosenbe@uta.edu

[§] Department of Industrial and Manufacturing Systems Engineering, The University of Texas at Arlington Campus Box 19017 Arlington, TX 76019 USA

2003 Pritsker Doctoral Dissertation award. Prior to joining the faculty at UTA, Dr. Rosenberger worked in the Operations Research and Decision Support Department at American Airlines.

Victoria Chen *is an associate professor of industrial engineering at The University of Texas at Arlington. She holds a PhD in Operations Research and Industrial Engineering from Cornell University. Dr. Chen's primary research interests utilize statistical methodologies to create new methods for operations research problems appearing in engineering and science. She has expertise in the design of experiments and statistical modeling, particularly for computer experiments and stochastic optimization. She has studied applications in inventory forecasting, airline optimization, water reservoir networks, wastewater treatment, and air quality. Through her statistics-based approach, she has developed computationally-tractable methods for continuous state stochastic dynamic programming, yield management, and environmental decision-making.*

Refined Experimental Design and Regression Splines Method for Network Revenue Management

Key Words: revenue management, Markov decision process, orthogonal array, multivariate adaptive regression splines

Abstract

We present a refinement of a network revenue management method that employs design of experiments and multivariate adaptive regression splines to approximate upper and lower bounds for the Markov decision process (MDP) value function. This approach involves an off-line *statistical modeling module* that approximates the value function to provide a policy for accepting/rejecting customer booking requests, and an on-line *availability processor module* that conducts the actual decisions as the booking requests arrive. In the statistical modeling module, the data for the value function upper and lower bound functions are obtained by solving deterministic and stochastic linear programming problems, respectively. The refinement in this paper identifies realistic ranges of remaining seat capacity at different reading periods by adding a *state space simulation module* preceding the statistical modeling module. This effectively combines the advantages of a design of experiments approach and a simulation-based approach. Simulation results using a real airline network and based on actual demand data demonstrate up to 2.7% improvement over the original method before refinement, which corresponds to a 5.8% improvement over the bid price approach using the deterministic linear programming model to determine bid prices. The state space simulation module could also be applied to improve approximate dynamic programming methods that conduct value function approximation.

1 Introduction

Revenue management (RM), also known as yield management, is defined as, “Selling the right seat at the right time to the right passenger for the right price” (Vinod 1995). RM is applied in various transportation sectors, such as auto rentals, ferries, rail, tour operators, cargo, and cruises. Other areas, like hotel/resorts, extended stay hotel, health care, manufacturing apparel, and companies that produce perishable goods etc., also use RM (Bodily and Pfeifer 1973). There is competition among the airline carriers to expand and explore RM to improve their revenue. American Airlines, for example, reported an increase in revenue of 5% due to improved RM methods in 1992, which translated to $1.4 billion over a 3-year period (Smith et al. 1992).

The network RM method in this paper is based on the approach of Chen et al. (2003) that employs an orthogonal array (OA) experimental design and multivariate adaptive regression splines (MARS) to approximate upper and lower bounds for the RM Markov decision process (MDP) value function. They formulated the RM model as an MDP, similar to that of Lautenbacher and Stidham (1999). Traditionally, MDP solutions are obtained using dynamic programming (DP), which is computationally intensive. Hence, the approach of Chen et al. (2003) was motivated by the OA/MARS DP value function approximation method of Chen et al. (1999). In this present paper, we seek to improve upon the OA/MARS network RM method by refining the state space to a realistic range. Intuitively, a reduced and more realistic state space should enable more accurate function approximations; however, identifying these ranges is not trivial and has not been previously attempted.

Comprehensive RM overviews have been written by McGill and van Ryzin (1999), Talluri and van Ryzin (2004), and Phillips (2005). In practice, the most effective RM approaches are based on bid-pricing (Simpson 1989, Williamson 1992, Talluri and van Ryzin 1998). Bid pricing approaches approximate the "fair market value" of a seat with the sum the bid prices on the requested itinerary, where bid prices are obtained by solving for the dual of a revenue optimization formulation. These bid prices are typically updated at the reading dates. A booking request is accepted if the customer's bid is greater than the (estimated) fair market value of the requested itinerary. DP approaches based on the MDP formulation should yield more precise solutions than bid-pricing. The approach by Chen et al. (2003), on which this present paper is based, is most closely related recent approximate DP approaches (Bertsimas and Popescu 2003, Adelman 2006, Zhang and Cooper 2005, Bertsimas and de Boer 2005, and van Ryzin and Vulcano 2006). These will be discussed in Section 2, following a description of the MDP formulation and the OA/MARS network RM method of Chen et al. (2003). Section 3 presents our realistic state space generation and details on our implementation of the OA/MARS approach. Our computational results on a real airline network are presented in Section 4, and concluding remarks are given in Section 5.

2 Background and Motivation

A *leg* is a flight that travels non-stop from an origin to a destination. In an airline reservation system, customers request a particular *itinerary*, which consists of one or more legs. Prior to the booking period, an airline defines restrictions for each fare class, as described previously. In this paper, we refer to the combination of an itinerary and a fare class as a *demand class*. For each demand class, the airline determines a fare and a maximum number of seats allowed to be sold, called the *authorization level*. The booking process typically starts three months prior to the date of departure. In reality a customer makes a *booking request* that includes an itinerary and several other characteristics, such as the number of days in advance of the itinerary or whether the ticket is refundable. The airline assigns the customer to a demand class that has the same itinerary, has not already sold more seats than its authorization level, has characteristics consistent with those of the booking request, and has the lowest fare. The airline then offers to sell the customer a ticket for the requested itinerary at a price equal to the fare of the assigned demand class. Finally, the customer decides to accept or reject the offer. Authorization levels are often recalculated on certain specific dates during the booking process called *reading dates*. These dates get closer to each other as the day of departure gets closer. After the process, the airline collects data and estimates the demand for each demand class based upon the number of tickets purchased.

In our simulation, we randomly sample customers from the estimated demand of the demand classes. A customer makes a booking request by bidding a price equal to the fare of the demand class for the associated itinerary. Once the request is placed, an airline representative uses a computer reservation system to decide if the request is to be accepted or rejected. The customer's requested price is compared with a threshold calculated by the airline that represents the fair market value. If the customer's bid is higher than the airline's threshold value, then the request is accepted; otherwise, the request is rejected. Like the real booking process, the simulation reoptimizes the fair market value on reading dates. Booking requests rejected for any demand class due to unavailability or filled seats are called *spilled* demand. The major drawback to models like our simulation and others like it in the literature (e.g., Smith et al. 1992, Chen et al. 2003) is that demand for the demand classes is assumed to be independent. In reality, however, demand may be highly correlated, some of it may be recaptured with even high-priced demand classes, while some may be lost or gained due to pricing of other competitive airlines. Nonetheless, the independence assumption is commonly used in practice and in the literature, Moreover, its usage here does not compromise the merit of using statistical models to approximate the MDP value function, nor the approach to obtain the realistic state space described in this paper.

The state of the system consists of the remaining seat capacities on n flight legs in the network. There are T reading dates, defining *T reading periods*, where the t-th reading period begins at the t-th reading date furthest from departure. Denote the *state vector* at the beginning of reading period t by $\boldsymbol{x}_t = (x_{1t}, \ldots, x_{nt})$. The *decision vector* in reading period t, denoted by $\boldsymbol{u}_t = (u_{1t}, \ldots, u_{mt})$, consists of the seat allocations for m demand classes. Let $A = (a_{if})$ denote the *leg-itinerary incidence matrix* with $a_{ij} = 1$ if leg i is utilized in demand class f and 0 otherwise. The *revenue* associated with demand class f is denoted by r_f (vector $\boldsymbol{r}$). The *total seat demand* for demand class f realized in reading period t is denoted by the random variable d_{ft}. Finally, the *value function* $F_t(\boldsymbol{x}_t)$ provides the maximum expected revenue accrued from reading periods t to departure. In Section 2.1, the MDP value function is described, and its approximation via the OA/MARS network RM method is given in Section 2.2. Section 2.3 discusses related approximate DP methods and provides motivation for both the OA/MARS approach and the refinement in this paper.

2.1 RM as an MDP

In the airline booking process, the decision to accept or reject a current booking request depends on the remaining seat capacity, the time the request was placed, the itinerary and fare class requested, and other characteristics of the current request, but it does not depend on decisions made about previous booking requests. Hence, RM can be classified as an MDP (Lautenbacher and Stidham 1999). The MDP formulation for the RM problem divides the three-month booking period into t_{MDP} time intervals, with at most one booking request per interval. These intervals are indexed in decreasing order, $k = t_{\text{MDP}}, \ldots, 1, 0$, where k = 1 denotes the first interval immediately preceding departure, and k = 0 is at departure. These intervals are different from the reading periods since each reading period can have multiple booking requests while each MDP interval can have at most one booking request.

Let $p_i^f(g)$ denote the probability that a request for g seats of demand class f occurs in time interval k; $p_k(0)$ denotes the probability of no booking requests in time interval k; and G_f is the maximum size of a group request for demand class f. Suppose the booking process is in state $\boldsymbol{x}$ at the beginning of time interval k. If a booking request for g seats that arrives during time interval k is accepted, then a new state $\boldsymbol{x}'$ is reached at the beginning of time interval $k-1$, where $\boldsymbol{x}'$ subtracts g seats from the legs involved in the requested itinerary and is greater than or equal to zero. The MDP value function $F_k(\boldsymbol{x})$, for $\boldsymbol{x} \geq 0$, provides the maximum expected revenue collected over time intervals k through departure when the system is at state $\boldsymbol{x}$ at the beginning of time interval k. Then $F_0(\boldsymbol{x}) = 0$ for all x. Thus, the MDP value functions can be written as:

$$F_k(\boldsymbol{x}) = \sum_{f=1}^{m} \sum_{g=1}^{G_f} p_k^f(g) \begin{cases} \max\left\{ g r_f + F_{k-1}(\boldsymbol{x}'), F_{k-1}(\boldsymbol{x}) \right\}, & \text{if } (\boldsymbol{x}' \geq 0) \\ F_{k-1}(\boldsymbol{x}) & \text{otherwise.} \end{cases}$$

If $\boldsymbol{x}' \geq 0$, then the airline will accept the booking request if $g r_f + F_{k-1}(\boldsymbol{x}') > F_{k-1}(\boldsymbol{x})$; otherwise, it will be rejected (Talluri and van Ryzin 2004). Hence, the *fair market value* (FMV) of a group of requested seats is defined as the difference in the value function of rejecting the request versus accepting the request: $\text{FMV} = F_{k-1}(\boldsymbol{x}) - F_{k-1}(\boldsymbol{x}')$ where $\boldsymbol{x}' \geq 0$. The RM policy is then to accept the booking request if the "bid" $g r_f$ is greater than FMV. As mentioned earlier, bid prices are used in practice to estimate FMV.

2.2 An OA/MARS Network RM Approach

Motivated by the successful application of design of experiments and regression splines in stochastic DP (Chen et al. 1999), Chen et al. (2003) proposed an OA/MARS network RM method. In this approach, the RM problem is solved in two parts, off-line and on-line. The off-line *statistical modeling module* derives the RM accept/reject policy while the on-line *availability processor module* conducts the actual decisions. Their model assumes:

1. The booking process starts ninety days before the day of departure.
2. Flight capacities and the flight schedule are known.
3. There is no overbooking or cancellation.

2.2.1 Bounding the MDP Value Function

Instead of directly approximating the value function, an OA/MARS approach is used to approximate upper and lower bounds. At reading date *t*, we have remaining capacities $\boldsymbol{x}_t$, and a deterministic (DET) linear programming problem may be solved to obtain the seat allocations $\boldsymbol{u}$ aggregated over reading dates *t* through departure, given the expected demand $E[\boldsymbol{d}]$ over reading dates *t* through departure:

$$\begin{aligned} \text{(DET)} \max \boldsymbol{ru} \\ \text{s. t. } A\boldsymbol{u} \le \boldsymbol{x}_t \\ 0 \le \boldsymbol{u} \le E[\boldsymbol{d}]. \end{aligned}$$

Talluri and van Ryzin (1998) showed that the maximum revenue from (DET) provides an upper bound on the value function $F_t(\boldsymbol{x}_t)$, and we denote this upper bound function by $F_t^U(\boldsymbol{x}_t)$. A stochastic model, also called the probabilistic nonlinear programming model is considered when demand is a random variable. This model is also referred to as the stochastic (STOCH) network model and is as given below:

$$\text{(STOCH)}\max \sum_{\tau=1}^{t}\sum_{f=1}^{m} r_f E[\min(d_{f\tau}, u_{f\tau})]$$

$$\text{s. t. } A\left(\sum_{\tau=1}^{t} \boldsymbol{u}_\tau\right) \le \boldsymbol{x}_t$$

$$u_{ft} \ge 0.$$

Günther (1998) showed that the maximum revenue from (STOCH) provides an lower bound on the value function $F_t(\boldsymbol{x}_t)$, and we denote this lower bound function by $F_t^L(\boldsymbol{x}_t)$. The duals of (DET) and (STOCH) are often used to provide bid prices for each flight leg (Williamson 1992, D'Sylvia 1982).

2.2.2 Statistical Modeling Module

The steps involved in this module are:

1. The reading dates are chosen and remaining seat capacities are initially set equal to the flight capacities for the flight legs.
2. An OA experimental design is constructed to provide discretized coverage of the remaining seat capacity state space. The state space ranges from zero to the plane capacities of the flights in the network.
3. For each of the discretization points, the (DET) model and the (STOCH) model are solved to provide points on the upper and lower bound functions $F_t^U(\boldsymbol{x}_t)$ and $F_t^L(\boldsymbol{x}_t)$, respectively. This loop is repeated at all reading dates.
4. For each reading date, a MARS approximation is fit separately to estimate $F_t^U(\boldsymbol{x}_t)$ and $F_t^L(\boldsymbol{x}_t)$ over the entire state space of $\boldsymbol{x}_t$. Thus, a total of $2T$ different MARS approximations are generated. We denote these by $\hat{F}_t^U(\boldsymbol{x}_t)$ and $\hat{F}_t^L(\boldsymbol{x}_t)$.

Figure 1 illustrates the procedure in the statistical modeling module. The MARS statistical models $\hat{F}_t^U(\boldsymbol{x})$ and $\hat{F}_t^L(\boldsymbol{x})$ are now available for the on-line availability processor module. For a 20-city, 31-leg network, computational effort per reading date takes less than 10 minutes on a 2.8-GHz Intel Xeon Workstation.

Figure 1 about here.

2.2.3 Availability Processor Module

At time τ between reading dates t and $t-1$, we use linear interpolation between $\hat{F}_t^U(\boldsymbol{x})$ and $\hat{F}_{t-1}^U(\boldsymbol{x})$ to approximate $F_\tau^U(\boldsymbol{x})$; similarly, for $F_\tau^L(\boldsymbol{x})$. An FMV for a booking request of group size *g*, for demand class *f* at time τ is estimated using

$$\text{Pessimistic} = \hat{F}_\tau^L(\boldsymbol{x}) - \hat{F}_\tau^U(\boldsymbol{x}')$$
$$\text{Optimistic} = \hat{F}_\tau^U(\boldsymbol{x}) - \hat{F}_\tau^L(\boldsymbol{x}')$$
$$\text{Estimated Fair Market Value} = \frac{\text{Pessimistic} + \text{Optimistic}}{2}$$

In Figure 2, the RM policy is defined as, “accept the booking request only if the requested fare is greater than the estimated fair market value.”

Figure 2 about here.

2.3 Approximate DP and Motivation

2.3.1 Discussion of Approximate DP Methods in RM

Both Chen et al. (2003) and Bertsimas and Popescu (2003) use the (DET) model to approximate the MDP value function in the definition of FMV. The distinction is that Bertsimas and Popescu (2003) require solving (DET) every time FMV must be estimated, while Chen et al. (2003) used (DET) and (STOCH) to construct a MARS-based value function approximation off-line, then employed these quick-to-compute MARS approximations to estimate FMV during the on-line booking process.

Adelman (2006) incorporates the linear programming approach to approximate DP (see de Farias and Van Roy 2003), which models the value function as a linear combination of “features” or basis functions, where the difficult task is identifying the appropriate set of features. Adelman (2006) uses an affine functional approximation with the remaining capacities as the features. Further, Adelman (2006) derives theoretical connections between this approximate DP approach and the bid price approach. Computational times were reasonable: for a 100-day time horizon, times were less than 4 minutes for a 20-leg hub-and spoke network and less than 157 minutes for a 40-leg hub-and spoke network on a 3.6GHz Intel Xeon Workstation.

Zhang and Cooper (2005), Bertsimas and de Boer (2005), and van Ryzin and Vulcano (2006) all consider a simulation-based approaches. Zhang and Cooper (2005) derived upper and lower bounds for the MDP value function of a dynamic customer choice model and then solved this stochastic optimization problem via simulation-based approaches. Bertsimas and de Boer (2005) employed fine-grained simulation-based state and decision space discretization and multi-linear interpolation to conduct value approximation, which required 50-60 hours to compute for a 30-leg hub-and-spoke network. By contrast, van Ryzin and Vulcano (2006) focused on a stochastic approximation approach with a continuous state and decision model to derive a computationally faster continuous optimization problem that overcomes the computational burden of the approach by Bertsimas and de Boer (2006). Computational times were approximately a minute per reading date for a 62-leg network.

Although a stochastic approximation approach is faster than a value function approximation approach, the value function approximation approach can yield better solutions (Wen 2005).

The critical issue in speeding up a value function approximation approach is the state space discretization. The novelty of the OA/MARS DP method of Chen et al. (1999) and the OA/MARS network RM method of Chen et al. (2003) is the use of *design of experiments* to obtain an efficient state space discretization. We discuss the motivation for this approach next.

2.3.2 Motivation for Refined OA/MARS Network RM Method

Design of experiments dates back to Fisher (1925), Bose and Kishen (1940), and Cochran and Cox (1950), when the necessity of minimizing the number of experimental runs to achieve a desired modeling precision was dictated by the cost of physical experiments. More recently, computer experiments have become a popular method to study complex systems. Although the cost of physical experiments is not an issue, design and analysis of computer experiments evolved to efficiently conduct computationally expensive computer experiment runs (see Sacks et al. 1989, Koehler and Owen 1996, Chen et al. 2006).

RM methods that seek to approximate a value function via a simulation-based approach automatically generate sample paths that represent a realistic state space; however, they do not take advantage of the efficiencies gained via design of experiments. A sample path that realizes states close to previously realized states provides little new information in approximating a response surface of a computer experiment, such as a value function. A simulation-based approach cannot control this, but design of experiments can, by spreading the discretization points over the state space, so that each point provides valuable information towards the approximation of the response surface. *The approach in this paper is the first to combine the benefits of a realistic state space and design of experiments.*

The statistical modeling method is also a critical choice, as that determines how well the given computer experiment data (e.g., points solved on the value function) can be employed to produce an accurate, yet, parsimonious approximation. Splines were shown to be much more efficient at using data than multi-linear interpolation for value function approximation (Johnson et al. 1993), and MARS combined with the efficiency of OA experimental designs can break the exponential growth referred to in DP as the curse of dimensionality (Chen et al. 1999). Neural networks have also been studied for this purpose (e.g., Bertsekas and Tsitsiklis 1996, Cervellera et al. 2006) and can provide comparable accuracy to MARS; however, there are complications in determining neural network architecture while MARS model structure is adaptively chosen by the MARS algorithm to fit the given data. Finally, it should be noted that neither the affine functional approximation of Adelman (2006) nor bid price approaches incorporate possible interactions between the flight legs. Interactions occur between flight legs on popular itineraries with a connection. The MARS approximations for the 31-leg network of Chen et al. (2003) do involve several interactions terms.

3 Refined OA/MARS Network RM Method

In the refined version of the OA/MARS network RM method, realistic ranges of the remaining capacity state variable are generated, instead of the same ranges (from zero to capacity) throughout the booking period. Intuitively, these ranges should be close to the capacity at the beginning of the booking period and move closer to zero towards departure. However, other than this intuitive information, realistic ranges are unknown a prior. Hence, we achieve this by adding a *state space simulation module* preceding the statistical modeling module of Section 2.2.

3.1 State Space Simulation Module

In the statistical modeling module of Section 2.2.2, the state space remains the same for all the reading dates. As a consequence, the design points are spread out over a wider region than required. We know that, in practice, one is unlikely to find an empty flight on the day of the departure or a full flight 90 days prior to the day of departure. In order to be more realistic, we estimate the possible/realistic ranges for each reading date. These are called trust regions.

Demand scenarios are generated based on real data. Remaining flight capacity is initialized to the actual flight capacity. The (DET) model, as described in Section 2.2.1, is employed at the reading dates to generate bid prices. The RM policy which states, “accept the booking request only if the fare is greater than the bid price” is used to make decisions on accepting/rejecting the request. Upon accepting the request, remaining seat capacity is updated to remaining capacity minus the booking request’s group size *g*. Demand scenarios are simulated many times and at the end of each reading date, remaining seat capacities are recorded. Figure 3 shows the generation of the trust regions. Remaining seat capacities obtained at each reading date over the entire simulation are used to determine the maximum and minimum capacities at those reading dates.

Figure 3 about here.

To estimate the number of simulation runs needed to obtain good realistic ranges, a simulation for an initial sample size of s = 30 was run. The resulting data was used to estimate the standard deviation of remaining capacity, σ. Desired sample size was then estimated using a confidence interval approach,

$$s = \left\lceil \left(\frac{2 z_{\alpha/2} \sigma}{E} \right)^2 \right\rceil$$

where E is the desired width of the confidence interval and was chosen to be 5% of the expected value of the sample plus or minus the confidence coefficient times the standard error. (Note: If the determined sample size is less than 30, then a confidence interval based on a *t*-distribution could provide a more accurate calculation.)

3.2 Approximation of the Value Functions

The remaining seat capacity state spaces are set according to the empirically-derived realistic ranges from the state space simulation module. Similar to step 2 of the statistical modeling module in Section 2.2.2, an OA experimental design is employed to identify discretization points in the realistic state spaces for each reading date. Otherwise, steps 3 and 4 in Section 2.2.2 remain essentially the same (see Figure 1).

As in Section 2.2.1, the (DET) model is used to provide an upper bound on the MDP value function, and the (STOCH) model is used to provide a lower bound. Solving the deterministic model is a straightforward LP, but there are different approaches for solving the stochastic model. The next section describes the approach used in this paper.

3.3 Solving the Stochastic Network Optimization Model

In the (STOCH) model, let variable $w_{f\tau}$ be the expected number of passengers for demand class f at reading date t $\forall f = 1,2,...,m, \forall \tau = 0,1,...,t.$ We can impose this definition using the following constraint set:

$$w_{f\tau} \leq E[\min(d_{f\tau}, u_{f\tau})], \forall f = 1,2,...,m, \forall \tau = 0,1,...,t. \quad (1)$$

We can then replace the objective function by the following:

$$\max \sum_{\tau=1}^{t} \sum_{f=1}^{m} r_f w_{f\tau}.$$

For each $f = 1,2,...,m$ and each $\tau = 0,1,...,t$, the function $E[\min(d_{f\tau}, u_{f\tau})]$ is known to be concave, as long as the random variable $d_{f\tau}$ is continuous. In this case, it is well-known within convex programming (Kelley 1960) that constraint set (1) can be replaced by an infinite set of first-order, so (STOCH) is equivalent to following reformulation:

$$\max \sum_{\tau=1}^{t}\sum_{f=1}^{m} r_f w_{f\tau}$$

$$\text{s.t. } A\left(\sum_{\tau=1}^{t} u_\tau\right) \le x_t$$

$$w_{f\tau} \le E[\min(d_{ft}, u_0)] + (u_{f\tau} - u_0)\nabla E[\min(d_{f\tau}, u_0)]$$

$$\forall f = 1,2,...,m, \forall\, \tau = 0,1,...,t, \forall u_0 \in \Re^+ \qquad (2)$$

$$u_{f\tau} \ge w_{f\tau} \ge 0, \forall f = 1,2,...,m, \forall\, \tau = 0,1,...,t.$$

However, when the random variable $d_{f\tau}$ is discrete, the function $E[\min(d_{f\tau}, u_{f\tau})]$ is not smooth, so constraint set (2) is undefined. In particular, suppose demand follows a compound Poisson process with arrival rate λ. From the definition of expected value we know that, for any discrete random variable *L*, $E[L] = \sum lp(l)$. For simplifying the derivation, we omit the subscripts *f* and τ, and let $d_{f\tau} = d$ and $u_{f\tau} = u$. Hence,

$$E[\min(d,u)] = \sum_{d=0}^{u} dp(d) + \sum_{d=u+1}^{\infty} up(d),$$

where *p(d)* is the probability of demand *d*. Let *H* be the cumulative distribution function for the Poisson distribution and *h* be the probability mass function for the Poisson distribution. Hence, if $u \ge 1$, then

$$\begin{aligned}
E[\min(d,u)] &= \sum_{d=0}^{\lfloor u \rfloor} de^{-\lambda}\lambda^d / d! + u \sum_{d=\lfloor u \rfloor+1}^{\infty} e^{-\lambda}\lambda^d / d! \\
&= \lambda \sum_{d=1}^{\lfloor u \rfloor} e^{-\lambda}\lambda^{d-1} / (d-1)! + u \sum_{d=\lfloor u \rfloor+1}^{\infty} e^{-\lambda}\lambda^d / d! \\
&= \lambda \sum_{d=0}^{\lfloor u \rfloor-1} e^{-\lambda}\lambda^d / d! + u\left[1 - \sum_{d=0}^{\lfloor u \rfloor} e^{-\lambda}\lambda^d / d!\right] \\
&= \lambda H(u-1) + u\left[1 - H(u)\right].
\end{aligned}$$

For $1 > u \ge 0$, $E[\min(d,u)]$ is just equal to $u\left(1 - H(0)\right)$. Since the cumulative distribution function *H* is a step function, $E[\min(d,u)]$ is piecewise linear. We can show that it is also continuous and hence concave. Observe that for a sufficiently small increase in the value $u < \lceil u \rceil$, the slope of the function $E[\min(d,u)]$ is equal to $1 - H(\lfloor u \rfloor)$. Now consider the difference between $E\left[\min\left(d, \lfloor u \rfloor + 1\right)\right]$ and $E\left[\min\left(d, \lfloor u \rfloor\right)\right]$,

$$\begin{aligned}
E\left[\min\left(d,\lfloor u \rfloor+1\right)\right]-E\left[\min\left(d,\lfloor u \rfloor\right)\right] &= \left[\left(\lfloor u \rfloor+1\right)-\left(\lfloor u \rfloor+1\right)H\left(\lfloor u \rfloor+1\right)+\lambda H\left(\lfloor u \rfloor\right)\right] \\
&\quad -\left[\lfloor u \rfloor-\lfloor u \rfloor H\left(\lfloor u \rfloor\right)+\lambda H\left(\lfloor u \rfloor-1\right)\right] \\
&= \lambda h\left(\lfloor u \rfloor\right)-\lfloor u \rfloor h\left(\lfloor u \rfloor+1\right)-H\left(\lfloor u \rfloor+1\right)+1 \\
&= \lambda h\left(\lfloor u \rfloor\right)-\left(\lfloor u \rfloor+1\right)h\left(\lfloor u \rfloor+1\right)-H\left(\lfloor u \rfloor\right)+1 \\
&= e^{-\lambda}\lambda^{\lfloor u \rfloor+1}\big/\lfloor u \rfloor!-\left(\lfloor u \rfloor+1\right)e^{-\lambda}\lambda^{\lfloor u \rfloor+1}\big/\left(\lfloor u \rfloor+1\right)!+1-H\left(\lfloor u \rfloor\right) \\
&= 1-H\left(\lfloor u \rfloor\right).
\end{aligned}$$

Since the slope of the function $E[\min(d,u)]$ for a sufficiently small increase in $u < \lceil u \rceil$ is equal to the difference between $E\left[\min\left(d, \lfloor u \rfloor + 1\right)\right]$ and $E\left[\min\left(d, \lfloor u \rfloor\right)\right]$, the function $E[\min(d,u)]$ is continuous. Consequently, we can replace the constraints in set (2) by an infinite set of constraints associated with discrete integer values of u_0. Thus, the original (STOCH) model is equivalent to the following reformulation:

$$\begin{aligned}
&\max \sum_{\tau=1}^{t}\sum_{f=1}^{m} r_f w_{f\tau} \\
&\text{s.t. } A\left(\sum_{\tau=1}^{t} u_\tau\right) \le x_t \\
&w_{f\tau} - (1-H(u_0))u_{f\tau} \le E[\min(d_{ft},u_0)] - u_0(1-H(u_0)) \\
&\qquad \forall f = 1,2,\ldots,m, \forall \tau = 0,1,\ldots,t, \forall u_0 \in \mathbb{Z}^+ \qquad (3) \\
&u_{f\tau} \ge w_{f\tau} \ge 0, \forall f = 1,2,\ldots,m, \forall \tau = 0,1,\ldots,t.
\end{aligned}$$

The set (3) has an infinite number of constraints, so it is computational intractable to solve this linear programming problem exactly as stated. However, we can approximate this linear programming problem by replacing the infinite set $\mathbb{Z}^+$ by a finite set in which $u_0 = 1, 2, \ldots, \bar{u}$, where $\bar{u}$ in a practical upper bound on the values $u_{f\tau}$, for all $f = 1, \ldots, m$ and $\tau = 0, \ldots, t$. In our computational experiments in Section 5, we set ū to be 50.

4 Computational Results

Our methodology was tested on the 31-leg airline hub application used by Chen et al. (2003). Their application used fifteen reading dates, included 123 itineraries, and had a maximum demand of 50. Data on flight capacities and demand distribution parameters were provided by a domestic airline carrier.

Based on the confidence interval approach, a total of 85 simulation runs were conducted in the state space simulation module to obtain the realistic state spaces. The computational effort was less than 5 minutes on a 2.8-GHz Intel Xeon Workstation. An OA experimental design, was used to generate 31^2 design/discretization points. Demand scenarios were generated based on the data given. Using all this information the refined statistical modeling module was executed, and the resulting RM policy was employed in a simulated booking process. The results obtained from 2000 simulation runs are given in Table 1.

Table 1 about here.

The column labeled "Load" is the nominal load factor, defined to be the quotient of total requested capacity over available capacity. Airlines use nominal load factors of up to 150%. The column "CV" contains the coefficient of variations considered for each load factor. The average revenues generated in the simulation are reported for four RM methods: the (DET) bid price approach, the (STOCH) bid price approach, the original OA/MARS network RM method of Section 2.2, and the refined OA/MARS network RM method of Section 3. The standard error is given below each average revenue, followed by the percentage increase in average revenue compared with the (DET) bid price approach. The upper bound at each of the load factors are obtained by solving (DET) model at 90-days before the day of departure for maximum capacities.

5 Concluding Remarks

It can be seen in Table 1 that the original OA/MARS network RM method is better than both the (DET) and (STOCH) bid price approaches, and the refined original OA/MARS network RM method shows improvement over the original. We can also observe that the (STOCH) bid price approach performs better than the (DET) bid price approach in certain instances. These results not only demonstrate the benefit of an OA/MARS network RM method, but, further, the benefit of using realistic trust regions for the state space. Although the idea of employing a state space simulation module to define realistic trust regions was first tested here in this paper on the OA/MARS network RM method, its success suggests that the OA/MARS DP method and similar value function approximation approaches would also benefit.

In on-going research, we are relaxing the assumption of no overbooking or cancellation. It is possible to estimate the number of seats to be overbooked, called overbooking pad, for each flight leg before the booking process starts. In the statistical modeling approach, the state space would be modified to be the sum of the actual flight capacity and the overbooking pad. Using a MARS approximation of revenue, we are studying the use of a combined Newton’s and steepest ascent method to estimate the optimal overbooking pad.

References

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

Bertsekas, D. P. and Tsitsiklis, J. N.: 1996, *Neuro-Dynamic Programming*, Belmont, Massachusetts: Athena Scientific.

Bertsimas, D. and de Boer, S.: 2005, Simulation based booking limits for airline revenue management, *Operations Research* **53**(1), 90–106.

Bertsimas, D. and Popescu, I.: 2003, Revenue management in a dynamic network environment, *Transportation Science* **37**(3), 257–277.

Bodily, S. E. and Pfeifer, P. E.: 1973, Overbooking decision rules, *OMEGA* **20**, 129–133.

Bose, R. C. and Kishen, K.: 1940, On the problem of confounding in general symmetric factorial designs, *Sankhyā*, **5**, 21–26.

Cervellera, C., Chen, V. C. P. and Wen, A.: 2006, Neural network and regression spline value function approximations for stochastic dynamic programming, *Computers and Operations Research*, **34**, 70–90.

Chen, V. C. P., Günther, D. and Johnson, E. L.: 2003, Solving for an optimal airline yield management policy via statistical learning, *Journal of the Royal Statistical Society, Series C* **20**, 1–12.

Chen, V. C. P., Ruppert, D. and Shoemaker, C. A.: 1999, Applying experimental design and regression splines to high dimensional continues state stochastic dynamic programming., *Operations Research* **47**(1), 38–53.

Chen, V. C. P., Tsui, K.-L., Barton, R. R. and Meckesheimer, M.: 2006, Design, modeling, and applications of computer experiments, *IIE Transactions* **38**, 273–291.

Cochran, W. G. and Cox, G. M.; 1950, *Experimental Designs*, New York: Wiley.

D'Sylvia, E.: 1982, O and D seat assignment to maximize expected revenue, Technical report, Boeing Commercial Airplane Company, unpublished.

Fisher, R. A.: 1925, *Statistical Methods for Research Workers*, London: Oliver & Boyd.

Günther, D.: 1998, *Airline Yield Management*, Ph.D. thesis, Georgia Institute of Technology.

Johnson, S. A., Stedinger, J. R., Shoemaker, C. A., Li, Y. and Tejada-Guibert, J. A.: 1993, Numerical solution of continuous-state dynamic programs using linear and spline interpolation, *Operations Research*, **41**, 484–500.

Kelley, J. E.: 1960, The Cutting-Plane Method for Solving Convex Programs, *SIAM Journal on Applied Mathematics* **8**, 703-712.

Koehler, J. R. and Owen, A. B.: 1996, Computer experiments. In Ghosh, S. and Rao, C. R., editors, Handbook of Statistics, **13**, 261–308. New York: Elsevier Science.

Lautenbacher, J. C. and Stidham, S. J.: 1999, Underlying markov decision process in the single-leg airline yield-management problem, *Transportation Science* **33**(2), 136–146.

McGill, I. and van Ryzin, G. J.: 1999, Revenue management: Research overview and prospects, *Transportation Science*, **33**, 233–256.

Phillips, R.: 2005, *Pricing and Revenue Optimization*, Stanford University Press, Stanford, CA.

Sacks, J.,Welch,W. J., Mitchell, T. J. andWynn, H. P.: 1989, Design and analysis of computer experiments, *Statistical Science* **4**, 409–423.

Siddappa, S.: 2006, *Statistical Modeling Approach to Airline Revenue Management with Overbooking*, Ph.D. thesis, The University of Texas at Arlington.

Simpson, R. W.: 1989, Using network flow techniques for origin-destination seat inventory control, Technical Report, Massachusetts Institute of Technology, Flight Transportation Laboratory Memorandum, M89-1.

Smith, B. C., Leimkuhler, J. F. and Darrow, R. M.: 1992, Yield management at American Airlines, *Interfaces* **22**, 8–31.

Talluri, K. and Van Ryzin, G. J.: 1998, An analysis of bid- price control for network revenue management, *Management Science* **44**, 1577–1593.

Talluri, K. and Van Ryzin, G. J.: 2004, *The Theory and Practice of Revenue Management*, Kluwer Academic Publishers, Norwell, Massachusetts.

van Ryzin, G. J. and Vulcano, G.: 2006, Simulation-based optimization of virtual nesting controls for network revenue management, *Operations Research*, to appear.

Vinod, B.: 1995, Origin and destination yield management, in D. Jenkins (ed.), *Handbook of Airline Economics*, The Aviation Weekly Group of the McGraw-Hill Companies, New York, pp. 459–468.

Wen, A.: 2005, *Statistics-Based Approaches to Stochastic Optimal Control Problems*, Ph.D. thesis, University of Texas at Arlington.

Williamson, E. L.: 1992, *Airline Network Seat Inventory Control: Methodologies and Revenue Impacts*, Ph.D. thesis, Massachusetts Institute of Technology.

Zhang, D. and Cooper, W. L.: 2005, Revenue management for parallel flights with customer-choice behavior, *Operations Research* **53**(3), 415–431.

Table Caption

Table 1: Average revenues from 2000 simulations of the 31-leg hub using four methods: DET = Deterministic Bid Price, STOCH = Stochastic Bid Price, STAT = Original OA/MARS Network RM Method, REF STAT = Refined OA/MARS Network RM Method. Results are shown for load factors of 75%, 120%, and 150%, and various coefficient of variation (CV) values. Standard errors are given in parentheses, and percent increase in average revenue from DET is shown.

Figure Caption

Figure 1: Flow chart representing the statistical modeling module.

Figure 2: Flow chart representing the availability processor module.

Figure 3: Flow chart representing the state space simulation module for generation of realistic state spaces.

Table 1: Average revenues from 2000 simulations of the 31-leg hub using four methods: DET = Deterministic Bid Price, STOCH = Stochastic Bid Price, STAT = Statistical Modeling Approach, REF STAT = Refined Statistical Modeling Approach. Results are shown for load factors of 75%, 120%, and 150%, and various coefficient of variation (CV) values. Standard errors are given in parentheses, and percent increase in average revenue from DET is shown.

Load (%)	CV	DET	STOCH	STAT	REF STAT	Upper Bound
75	0.4	790427.2 (586.00)	789450.1 (533.60) -0.12%	790536.7 (503.50) 0.01%	790975.7 (486.50) 0.07%	799594.1
75	0.56	718729.4 (680.40)	720331.5 (762.60) 0.22%	725763.9 (623.00) 0.82%	728458.5 (645.30) 1.35%	799594.1
75	0.63	662294.1 (769.20)	662522.5 (858.70) 0.03%	665756.5 (845.40) 0.52%	668229.7 (764.40) 0.89%	799594.1
120	0.32	807252.8 (586.00)	801192.2 (508.00) -0.75%	818754.6 (455.20) 2.70%	820075 (467.50) 2.86%	993560.8
120	0.45	779258.5 (603.20)	775292.9 (548.40) -0.51%	808538.4 (457.70) 3.76%	817417.1 (428.40) 4.90%	993560.8
120	0.6	699243 (702.60)	709387.4 (774.00) 1.45%	720865.3 (736.80) 3.09%	739938.6 (632.70) 5.82%	993560.8
150	0.48	787231.3 (591.10)	784911.9 (500.60) -0.29%	790875.8 (524.60) 0.46%	799574.9 (425.65) 1.57%	1108892
150	0.56	766360.9 (609.40)	724377.4 (770.10) -6.90%	789653.5 (643.90) 3.04%	794364.7 (700.60) 3.65%	1108892
150	0.7	721903.3 (679.00)	607282.8 (779.30) -15.88%	749655.7 (756.90) 3.84%	753297.7 (737.80) 4.35%	1108892

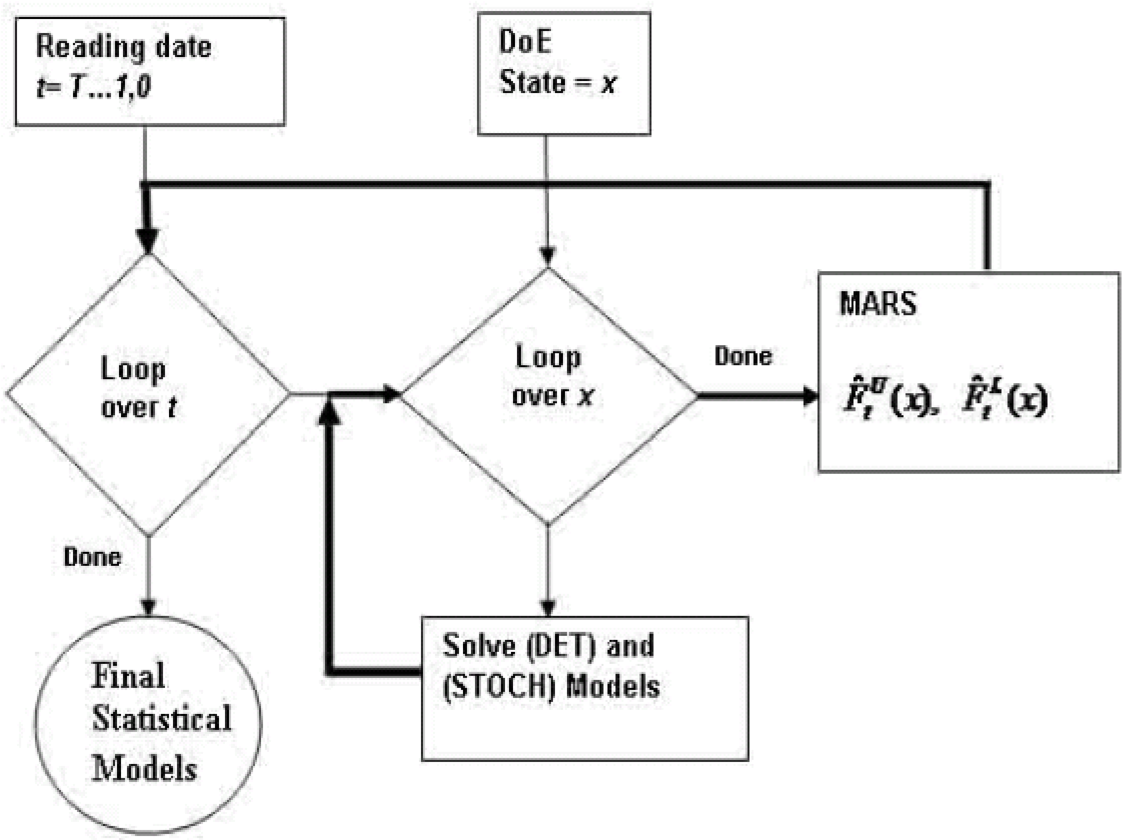

Figure 1: Flow chart representing the statistical modeling module.

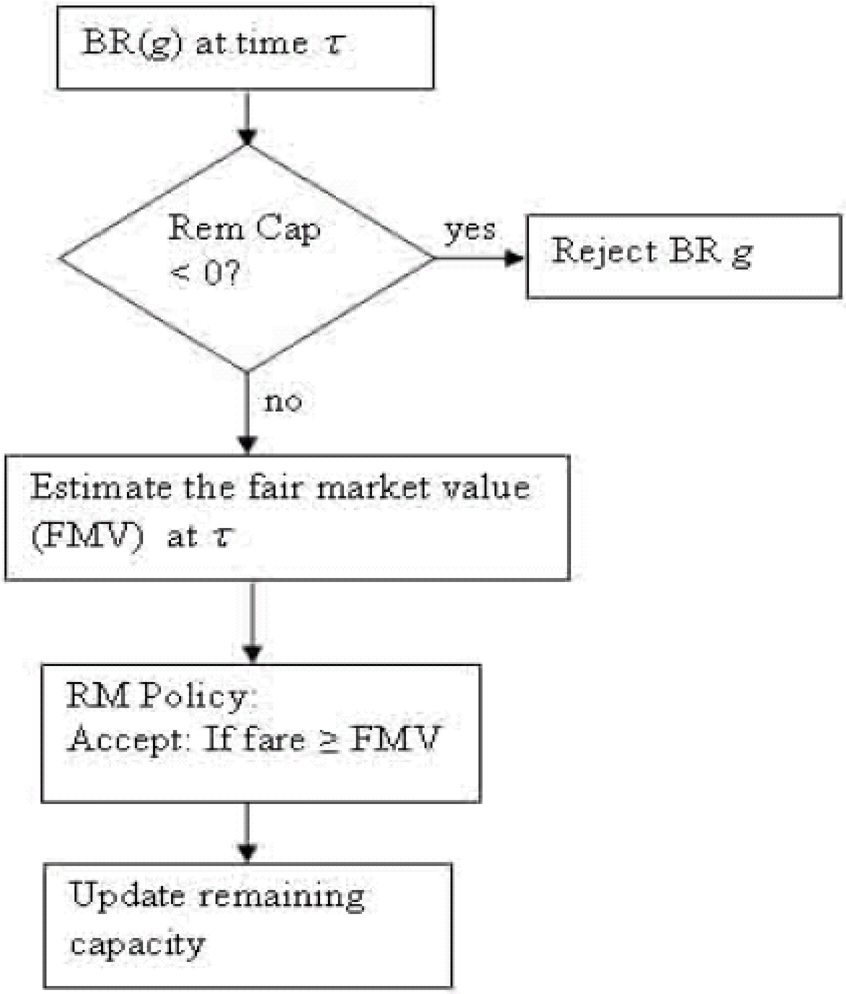

Figure 2: Flow chart representing the availability processor module.

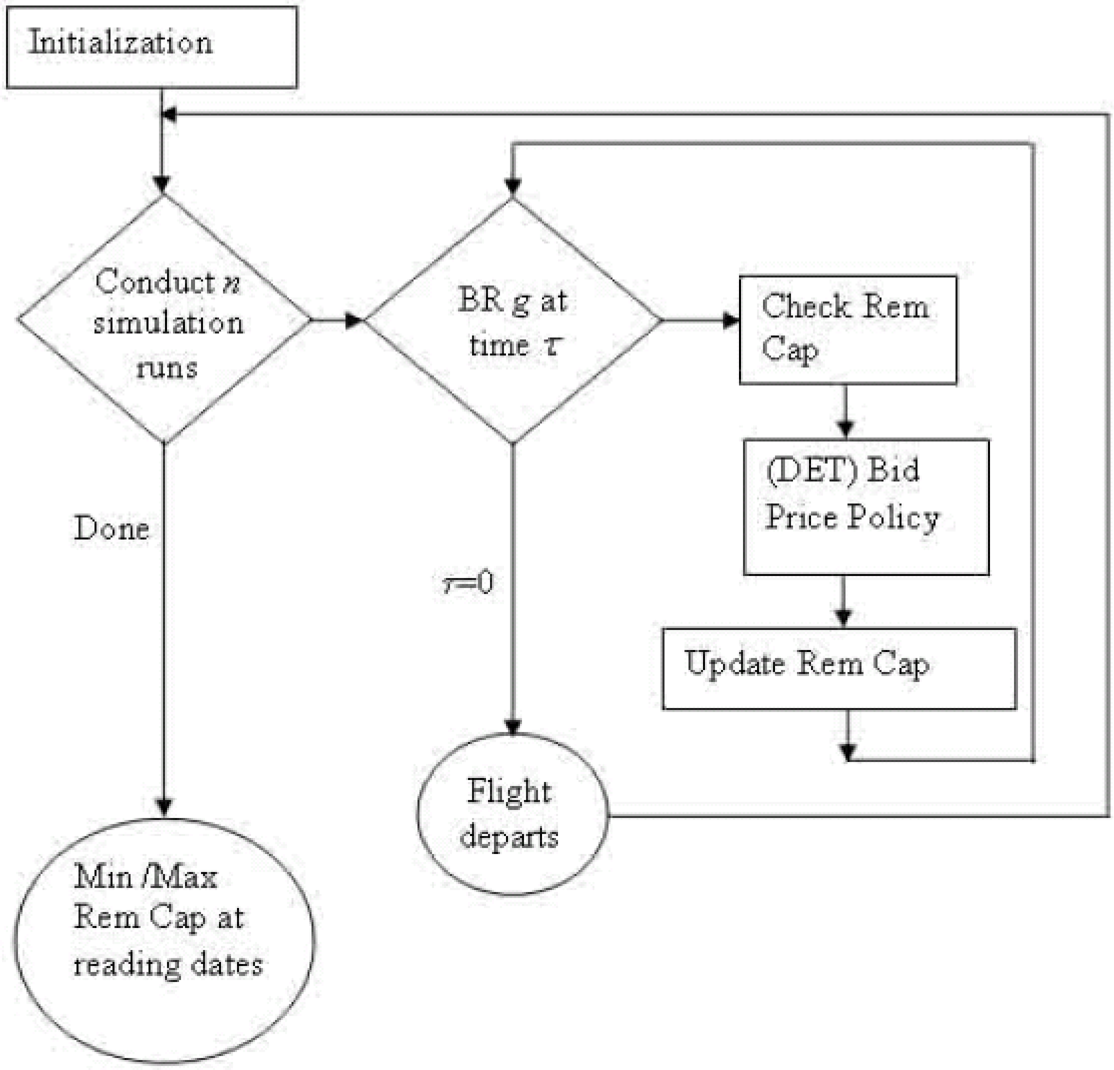

Figure 3: Flow chart representing the state space simulation module for generation of realistic state spaces.