

# Simulation Optimization for Revenue Management of Airlines with Cancellations and Overbooking

*Abhijit Gosavi*

Assistant Professor

Simulation-Optimization Lab

Department of Industrial Engineering

317 Bell Hall

SUNY, Buffalo

Buffalo, NY 14260

Email: agosavi@buffalo.edu, Tel: (716) 645-2357 EXT 2106

*Emrah Ozkaya*

Futek Advanced Sensor Technologies

Irvine CA 92618, USA

and

A.F. Kahraman

333 Bell Hall

SUNY, Buffalo, NY 14260.

November 16, 2006

### Abstract

This paper develops a model-free simulation-based optimization model to solve a seat-allocation problem arising in airlines. The model is designed to accommodate a number of *realistic* assumptions for real-world airline systems — in particular, allowing cancellations of tickets by passengers and overbooking of planes by carriers. The simulation-optimization model developed here can be used to solve both single-leg problems and multi-leg or network problems. A *model-free* simulation-optimization approach only requires a discrete-event simulator of the system along with a *numerical* optimization method such as a gradient-ascent technique or a meta-heuristic. In this sense, it is relatively “easy” because alternative models such as dynamic programming or *model-based* gradient-ascent usually require more mathematically-involved frameworks. Also, existing *simulation-based* approaches in the literature, unlike the one presented here, fail to capture the dynamics of cancellations and overbooking in their models. Empirical tests conducted with our approach demonstrate that it can produce robust solutions which provide revenue improvements over heuristics used in the industry, namely, EMSR (Expected Marginal Seat Revenue) for single-leg problems and DAVN (Displacement Adjusted Virtual Nesting) for networks.

## 1 Introduction

The field of airline revenue management studies *maximization* of revenues obtained by selling airline seats. An important problem in this field requires the development of a revenue-optimal strategy of customer selection. The product (airline seat) in airlines is said to have a “perishable” nature because its value becomes zero if it is not sold by the end of the booking horizon, which begins when the flight is opened for sale and ends when the flight takes off.

Typically, there are significant *differences* in the preferences (demands) of customers of an airline company. Some customers, usually business travelers, demand flexibility in cancellation options and return tickets within a week, while those traveling for leisure do not have these restrictions and opt for cheaper non-refundable tickets (with stiff cancellation penalties). Therefore, airline companies generally offer seats at *different* fares to *utilize* differences in passenger expectations *to their own advantage*. The number of business travelers is quite small in proportion, and business tickets are booked at the last minute, thereby making it important for the company to *retain* a few seats until the end of the booking horizon. The question that then arises is: how many seats should be allowed to be sold at any given fare? If one reserves too many seats for high-revenue passengers, it is possible that the plane will fly with many empty seats; on the other hand if all seats are sold at discount fares, one will potentially lose high-revenue passengers. Thus, an important

task is to *determine* the upper limit, called the *booking limit*, on the number of seats to be sold at or *allocated* to each fare offered.

The above-described “seat-allocation” problem is complicated by uncertainties in customer behavior and forecasts. Forecasts are generally prepared to estimate the probability distribution of the number of arrivals in each *fare class*. Inevitably, some passengers cancel tickets. Hence, airlines *overbook* planes in order to minimize the probability of flying with empty (canceled) seats, which adds to the complexity of the problem because cancellations are random. Thus seat-allocation should account for random cancellations and the feature of overbooking. Some realistic features of actual airline systems include: (i) random customer arrivals for booking, (ii) random cancellations, (iii) change in arrival rates with time, and (iv) concurrent (non-sequential order) arrivals of passengers, i.e., arrivals do *not* follow any particular order such as low fare classes first etc.

This paper studies the use of a *model-free simulation-optimization model* to solve the seat-allocation problem in a near-optimal manner. There are at least two reasons that make this approach attractive. Firstly, simulation can easily accommodate realistic assumptions (such as cancellations and overbooking), which often render theoretical models intractable. Secondly, *model-free* simulation-optimization models do not require knowledge of the *internal structure* of the stochastic system; all they need is an estimated numerical value of the objective function at any given point in the solution space, and a discrete-event simulator can provide these values easily. Actually, much research in recent times has allowed the efficient combination of simulation with numerical optimization techniques, such as gradient-ascent or meta-heuristics, which paves the way for generating implementable solutions. A commonly-prescribed method for simulation optimization in continuous spaces is the gradient-ascent approach that uses finite differences of the revenue function to estimate the gradient. A major difficulty with this approach is that the number of simulations required per iteration grows proportionately with the number of decision variables. Simultaneous perturbation (SP), which is due to Spall (1992), is a relatively new technique for gradient-ascent. The remarkable feature of this technique is that its computational burden is *not* proportional to the number of decision variables, and despite this, it has been shown to be convergent under certain conditions. It has been already used in countless applications (Spall, 2003). Although ours is a problem of discrete optimization, we have used a continuous *approximation* for solution purposes. Continuous (fluid) approximations for solving discrete problems are very common in the literature; see e.g., van Ryzin and Vulcano (2003). In addition to SP, we also apply simulated annealing (SA) on the problem. SA is a well-known meta-heuristic for discrete optimization and is known to be convergent only in an asymptotic sense. But since it has been widely applied in the industry, it will be used as a standard benchmark method in our computational experiments.

Because the revenue function will be estimated via simulation in this paper, the issue of simulation-induced noise — that can corrupt the function value — will be analyzed.

**Airline Networks:** An airline network is composed of one or more *hubs* and *spokes*. A *leg* in a hub-and-spoke network is composed of two cities. When a customer travels from one city to another via other cities, the itinerary of the customer contains multiple legs. Consider Figure 1 which shows a small network of four cities in the US. Chicago serves as the *hub*, and cities such as Miami, Denver, and Boston serve as *spokes*. Customers flying from one spoke city to another are routed via the hub. Most companies have one or two major hubs. Thus, Miami-Chicago-Boston forms one *origin-destination* itinerary or simply one itinerary. In large networks, one finds *several hundred* itineraries, and for each itinerary, multiple fares may be available; each itinerary-fare combination is often described as a *product*. The problem of seat-allocation can be studied either at the leg level (on a leg-by-leg basis) or at the network level. At the leg level, each fare is referred to as a class, while in the network, each itinerary-fare combination is referred to as a product. At the leg level, the problem is one of finding the number of customers to be allowed for a given fare (class). In a network, the seat-allocation problem is to determine the number of customers to be accepted for any given product. Complexity in the network version of the problem arises from the fact that seat-allocation on one leg *affects* that on one or more of the other legs. As a result, ideally, the airline that operates in networks must solve the problem in its entirety, i.e., solve the network version. In recent times, a large number of point-to-point carriers have emerged, which do *not* have hubs, but offer direct services from one city to another. For such carriers, it is sufficient to solve the problem at the leg level.

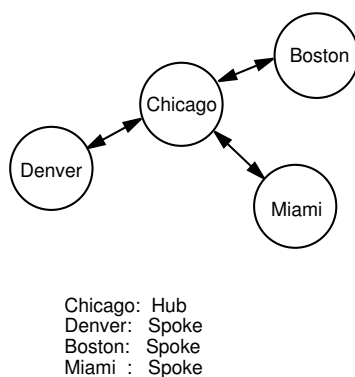


Figure 1: A schematic showing a network involving 4 cities

**A literature review and contributions of this work:** A sizable chunk of the work in single-leg revenue management exploits the equation underlying the pioneering work of Littlewood (1972), which has developed

into a very robust solution technique called EMSR (Expected Marginal Seat Revenue) described in Belobaba (1989). Displacement Adjusted Virtual Nesting (DAVN), which has its roots in the work of Glover et al. (1982), is a powerful approach for network models in revenue management. This model has been modified in subsequent years by several researchers including Smith and Penn (1988) and Williamson (1992). The literature on airline revenue management is quite voluminous. See McGill and van Ryzin (1999) for a review and Talluri and van Ryzin (2004) for textbook treatment of this topic, while the history is traced in Boyd and Bilegan (2003). Some important works in the domain of single-leg control are: Howard (1971); Shlifer and Vardi (1975); Curry (1990); Wollmer (1992); Lee and Hersh (1993); Brumelle and McGill (1993); Chatwin (1998); Robinson (1995). Subramaniam et al. (1999) present a finite-horizon Markov decision process to solve the single-leg problem, but make some limiting assumptions such as a Poisson rate for cancellations and equal cancellation probabilities for all classes. van Ryzin and McGill (2000) present a Robbins-Monro scheme that exploits simulation to solve the problem with forecasting as an integral part of the solution model; in most models in the literature, forecasts are assumed to be known. For the single-leg problem, Gosavi et al. (2002) present an approximate dynamic programming (DP) or reinforcement learning approach (Bertsekas and Tsitsiklis, 1996; Gosavi, 2003) that is based on value iteration and employs function approximation with neural networks for estimating the value function of DP within a simulator. Gosavi (2004) uses a policy iteration based algorithm in reinforcement learning for solving the same problem. Both of the above papers do not require the transition probabilities of the underlying stochastic dynamic program but work within simulators of airline systems.

For network control, outside of the pioneering paper of Glover et al. (1982), a subset of important works includes Vinod (1995); Smith and Penn (1988); Simpson (1989); Williamson (1992); Wong et al. (1993). Bertsimas and de Boer (2005) use a combination of gradient-ascent based on finite differences and approximate dynamic programming to solve the network problem. van Ryzin and Vulcano (2003) use a fluid approximation of the booking limits to obtain exact expressions for sub-gradients. The model in their paper exploits the structure of the problem, and can be viewed to belong to the class of *model-based* simulation-optimization algorithms (see Chapter 15 of Spall (2003)). Although both of the above papers use simulation, they are not designed to handle cancellations, which form an integral part of the booking process. Other approaches include: Experimental designs and multi-variate adaptive regression splines (Chen et al., 2003) in a DP setting, a sampling-based approach that combines merits of mathematical programming and Markov decision processes for a two-leg problem (Cooper and Homem-de-Mello, 2004), and an approximate DP algorithm for networks in which cancellations and overbookings are permitted (Bertsimas and Popescu, 2003). Karaesmen and van Ryzin (2004) is a recent paper that develops a cancellation-based model that exploits two-stage

stochastic programs. The work of de Boer et al. (2004) uses a combination of stochastic programming with simulation to derive booking limits for a network. Simulation optimization (Gosavi, 2003; Bonnans and Shapiro, 2000) is a rapidly growing area of research. The robust response surface method, which is a classical technique (Law and Kelton, 1999) for static simulation-based optimization, has made way for techniques that depend on meta-heuristics like tabu search (Glover, 1990), genetic algorithms (Holland, 1975), and simulated annealing (Kirkpatrick et al., 1983).

The contributions of our work, in the perspective of the existing literature, are (i) for the single-leg and network problems, we provide a *model-free* simulation-based optimization approach that can account for a variety of system-related assumptions, including arbitrary distributions for demand-arrival processes and cancellations, (ii) we introduce in the area of revenue management, the SP algorithm of Spall (1992), which other than in simulators (as we have done) could also be exploited in *model-based* gradient-ascent approaches in revenue management (van Ryzin and Vulcano, 2003; Bertsimas and de Boer, 2005), and (iii) we establish the usefulness of our simulation-optimization approach by analyzing the effect of simulation-induced noise in our computational experiments.

The rest of this paper is organized as follows. Section 2 presents a simulation-optimization model along with techniques used for solving the associated problem. Numerical results are discussed in Section 3. Section 4 concludes this paper.

## 2 A simulation-optimization approach

Two undesirable events are associated with *not* setting an upper limit on the number of seats to be sold at the different fares offered: 1) Passengers who end up in the plane at takeoff are *primarily* from the lower-revenue classes, which translates into loss of potential revenue. This is because lower-revenue passengers tend to book *early* and, if no booking limits are imposed, can buy all the seats in the plane. 2) Clearly, with no booking limits imposed, the number of passengers who show up for boarding can *exceed* the capacity of the plane. When this happens, some passengers have to be *bumped*, i.e., denied boarding request, although they have purchased tickets, which leads to loss of goodwill and revenue (arising from paying for tickets on alternative routes and hotel stays). However, if selling is discontinued as soon as the number of seats sold *equals* the capacity of the plane, some seats are likely to remain empty at takeoff — because of no-shows and last-minute cancellations.

We now introduce some mathematical notation. Each flight in a simulation model constitutes of arrivals,

cancellations, and the takeoff of the plane. Let  $\Omega$  denote the (universal) set of all possible flights. Consider a probability space  $(\Omega, \mathcal{F}, P)$  where  $\mathcal{F}$  is a sigma-field of subsets of  $\Omega$  and  $P$  denotes a probability measure on  $(\Omega, \mathcal{F})$ . In general, the seat-allocation problem can be described mathematically as follows. One has to find the values of the booking-limits vector  $\vec{x} = \{x_1, x_2, \dots, x_n\}$  that maximizes the expected revenue from each flight, i.e.,  $E_P [T(\vec{x}, \omega_i)]$ , where  $n$  denotes the number of classes (distinct fares) or products,  $E_P$  stands for the expectation operator induced by the probability measure  $P$ , and  $T : \mathbb{R}^n \times \Omega \rightarrow \mathbb{R}$  is a scalar-valued function. Here,  $T(\vec{x}, \omega_i)$  will denote the total (net) revenue obtained from the  $i$ th sample flight, denoted by  $\omega_i \in \Omega$ , in which the booking-limits vector is fixed at  $\vec{x}$ . In order to explain the working mechanism of a “simulation-optimization” approach, we need to make the problem statement more precise. A typical stochastic optimization problem is of the form

$$\max_{\vec{x} \in \theta} f(\vec{x}) \equiv \max_{\vec{x} \in \theta} E_P [T(\vec{x}, \omega_i)] \quad (1)$$

where  $\theta$  is a compact subset of  $\mathbb{R}^n$  and  $\vec{x}$  denotes the vector of decision variables. The random variables will have  $\omega$  in the notation to distinguish them from other quantities. Using the distribution  $P$ , if  $N$  i.i.d random samples,  $\omega_1, \omega_2, \dots, \omega_N$ , are generated, we can develop an approximation of the stochastic optimization problem in (1). The approximation is

$$\max_{\vec{x} \in \theta} \sum_{i=1}^N \frac{T(\vec{x}, \omega_i)}{N}, \quad (2)$$

where

$$f(\vec{x}) \equiv E_P [T(\vec{x}, \omega_i)] \approx \sum_{i=1}^N \frac{T(\vec{x}, \omega_i)}{N} = f_N(\vec{x}). \quad (3)$$

Now, (2) can be used in simulation-based optimization using a “sufficiently” large value for  $N$ . Although (2) is only an approximation of the problem in (1), it can be shown from the strong law of large numbers that as  $N \rightarrow \infty$ , (2)  $\rightarrow$  (1) with probability 1 (almost surely). Also, we will assume that if  $\lim_{N \rightarrow \infty} \frac{\sum_{i=1}^N T(\vec{x}, \omega_i)}{N}$  does not exist, we will maximize  $\liminf_{N \rightarrow \infty} \frac{\sum_{i=1}^N T(\vec{x}, \omega_i)}{N}$ .

The problem then is to formulate a procedure to estimate  $T(.,.)$ . The estimation can be performed in a simulator provided one formulates a suitable expression for  $T(.,.)$  — a task that we accomplish in the next subsection.

## 2.1 A simulation model

We  $i$ th aggregate sample of flights over all legs will be defined as:  $\hat{\omega}_i = \{\omega_i^1, \omega_i^2, \dots, \omega_i^l, \dots, \omega_i^L\}$ , where  $L$  denotes the number of legs in the network and  $\omega_i^l$  denotes the  $l$ th leg of the  $i$ th aggregate sample of flights.

Also, we define the following terms.  $H$ : the booking time horizon (in days);  $n$ : the number of products;  $M_v$ : penalty incurred by a passenger for cancellation of the  $v$ th product;  $V_v$ : revenue associated with the  $v$ th product;  $B_v$ : penalty incurred by the company for bumping a passenger of the  $v$ th product;  $C^l$ : plane's capacity in the  $l$ th leg;  $\bar{\rho}_v(\vec{x}, \hat{\omega}_i)$ : the number of passengers admitted for the  $v$ th product in the network if the booking-limits vector is  $\vec{x}$ ;  $\rho_v^l(\vec{x}, \omega_i^l)$ : the number of passengers admitted for the  $v$ th product in the  $l$ th leg if the booking-limits vector is  $\vec{x}$  (clearly, this quantity will always equal 0 if the  $v$ th product does not use the  $l$ th leg);  $\bar{c}_v(\vec{x}, \hat{\omega}_i)$ : the number of passengers who cancelled tickets for the  $v$ th product in the network if the booking-limits vector is  $\vec{x}$ ;  $c_v^l(\vec{x}, \omega_i^l)$ : the number of passengers who cancelled tickets for the  $v$ th product in the  $l$ th leg if the booking-limits vector is  $\vec{x}$  (clearly this quantity will also be 0 if the  $v$ th product does not use the  $l$ th leg).

Then, the *gross revenue* obtained in the  $i$ th aggregate sample of the network from selling seats and from cancellations, with  $\vec{x}$  as the booking-limits vector, can be expressed as

$$G(\vec{x}, \hat{\omega}_i) = \sum_{v=1}^n (\bar{\rho}_v(\vec{x}, \hat{\omega}_i) - \bar{c}_v(\vec{x}, \hat{\omega}_i)) V_v + \sum_{v=1}^n \bar{c}_v(\vec{x}, \hat{\omega}_i) M_v.$$

Also, the number bumped in the  $l$ th leg of the  $i$ th aggregate sample will be

$$K^l(\vec{x}, \omega_i^l) = \max \left[ 0, \left\{ \sum_{v=1}^n (\rho_v^l(\vec{x}, \omega_i^l) - c_v^l(\vec{x}, \omega_i^l)) - C^l \right\} \right].$$

Let  $\bar{K}^v(\vec{x}, \hat{\omega}_i)$  denote the number bumped for the  $v$ th product in the  $i$ th aggregate sample, which can be determined from the values of

$K^1(\vec{x}, \omega_i^1), K^2(\vec{x}, \omega_i^2), \dots, K^L(\vec{x}, \omega_i^L)$  and the simulator. Then, the net network revenue, associated with the  $i$ th sample, can be calculated from:

$$T(\vec{x}, \hat{\omega}_i) = G(\vec{x}, \hat{\omega}_i) - \sum_{v=1}^n B_v \bar{K}^v(\vec{x}, \hat{\omega}_i).$$

In our simulation-optimization approach, a booking-limits vector of  $\vec{x}$  implies that a customer requesting the  $i$ th product is accepted only if the net number (accepted number minus the cancellations) of customers currently booked for all products including  $i$  and those below  $i$  is less than  $x_i$ . In the single-leg scenario, the products are ranked by fares, and in the network by their net worth to the network, which is explained later.

Sample flights can be simulated in a computer program, and thereby one can compute the function  $T(.,.)$  for  $N$  sample flights using the definition above. Then (3) can be used to estimate the value of the objective function.

We now describe in detail the two techniques that we have used for simulation optimization.

## 2.2 Simultaneous Perturbation

SP, as mentioned previously, is an efficient steepest-ascent technique that can be used to solve continuous optimization problems with a large number of decision variables and a noisy function evaluator, e.g., a simulator. It is particularly suitable for simulation optimization Spall (2003). The gradient estimate of SP, unlike traditional finite differences, requires only two function estimates — that is, two simulations. This is the primary advantage of this method, and makes it suitable for our problem in which we have many decision variables. For a problem with  $n$  decision variables, finite difference approaches on the same problem would take  $n$  times as many simulations *per iteration* of the steepest-ascent algorithm.

We must note that SP is guaranteed only local convergence unless the function is concave or unimodal. It is extremely unlikely that the function in our problem domain is concave (even quasi-concave (van Ryzin and Vulcano, 2003)). As a result, in practice, one can start the search at a number of different points and select the best of the local optima as the solution. The steps, described in detail next, are written in terms of *maximizing* the objective-function value.

*Step 1.* Set  $k = 1$  and  $\bar{x}^1 = \{x_1^1, x_2^1, \dots, x_n^1\}$  to an arbitrary feasible solution in  $\mathcal{X}$ , the feasible set of the solution space, which as a rule of thumb can be set to  $\{0, 1, 2, \dots, M\}$ , where  $M$  is the maximum number of customers that can arrive in the booking horizon. Denote  $\mathcal{S} = \{1, 2, \dots, n\}$ . Set the step size  $\mu$  to a small value  $\mu^{small}$ . Set  $\mu_{min}$  to a small value such as 0.001.

*Step 2.* A random value for each  $J(i)$  is generated by simulation, where  $i \in \mathcal{S}$ , from the Bernoulli distribution whose two *equally-likely* values are 1 and  $-1$ . Thus  $J(i)$  is either 1 or  $-1$  with the same probability. We set  $d = 1/k^w$  where  $k^w$  denotes  $k$  raised to the power  $w$ , and  $w$  is fixed to a value in  $(0, 1)$ , e.g., 0.5. Then the perturbation parameter,  $h_i$  for every  $i \in \mathcal{S}$  is computed as follows:  $h_i = J(i)d$ .

*Step 3.* Let  $f$  denote the noisy function value obtained from a simulator. Thus  $f$  denotes the estimate obtained in Equation (3) using  $N$  samples from the simulator. Calculate  $F^+$  and  $F^-$  using the following:

$$F^+ = f(x_1^k + h_1, x_2^k + h_2, \dots, x_n^k + h_n), \quad F^- = f(x_1^k - h_1, x_2^k - h_2, \dots, x_n^k - h_n).$$

*Step 4.* For each  $i$  in  $\mathcal{S}$ , first obtain the derivative estimate and then update  $x_i^k$  via steepest descent:

$$\left. \frac{\partial f(\bar{x})}{\partial x_i} \right|_{\bar{x}=\bar{x}^k} \simeq \frac{F^+ - F^-}{2h_i}, \quad x_i^{k+1} \leftarrow \Pi_{\mathcal{X}} \left[ x_i^k + \mu \left. \frac{\partial f(\bar{x})}{\partial x_i} \right|_{\bar{x}=\bar{x}^k} \right],$$

where  $\Pi_{\mathcal{X}}[\cdot]$  denotes the projection operator onto the feasible set  $\mathcal{X}$ .

*Step 5.* Increment  $k$  by 1, and set:  $\mu \leftarrow \frac{\mu^{small}}{k}$ . If  $\mu < \mu_{min}$ , stop. Otherwise, return to Step 2.

When the function is evaluated in Step 3, one must round off each element of the vector  $\vec{x}$  to the nearest integer. We need to discuss the effect of simulation-induced noise on the objective function value. Fortunately, the simulator has a regenerative structure in which a new i.i.d. sample is generated whenever the flight takes off. Also, the objective function in (3) is the expected net revenue in one flight, and the problem is one of a finite horizon. As a result, the question of computing profits on a unit-time basis, which can lead to an additional bias (see the excellent discussion on pages 374-379 of Spall (2003)), does not arise. In fact, the estimator  $f_N(\vec{x})$  is an *unbiased* estimator of  $f(\vec{x})$  provided i.i.d samples are generated for  $T(\cdot)$ , which we ensure in our simulator.

The SP algorithm is guaranteed to converge under a number of conditions that we enumerate below. The conditions related to the step-size can be easily met in our computational experiments. In fact the step-size rules that we stated in our algorithm description obey these conditions. But the conditions related to the function gradients, e.g., differentiability and concavity, are difficult to verify. Spall himself notes in his book (see page 161 of Spall (2003)) that verifying the convergence conditions when the gradient is computed numerically is an “abstract ideal” and that the conditions “may not be verifiable for the very reason” one is forced to use numerical gradient differences. Fu and Hill (1997) state: “In practice, it may be difficult to verify the conditions on the objective function, since simulation is applied to those systems for which analytical properties are not readily available.” Some of these conditions could be verified only when expressions for gradients could be derived — in which case these expressions could be directly used in the optimization — *making a numerical evaluation of the gradients unnecessary*. The objective function considered in this paper has a black-box nature — especially when realistic assumptions are made for the system, and this is precisely why we pursue a simulation-based approach.

The conditions and the result that we state next are based on work in Fu and Hill (1997), which applies in our setting, i.e., constrained optimization in a compact subset.

#### Step-size and variable-parameter conditions:

1. Let  $\mu_k$  denote the step-size in the  $k$ th iteration of the algorithm. Then:  $\sum_{k=1}^{\infty} \mu_k = \infty$ .
2. For every  $i$ , the sequence of random variables  $J(i)$  are mutually independent with mean 0, have bounded second moments, and  $E[1/J(i)]$  is uniformly bounded.
3. If  $e_k(i)$  denotes the noise in the gradient due to the use of SP, and  $\vec{e}_k = [e_k(1), e_k(2), \dots, e_k(n)]^T$ , then  $\sum_{k=1}^{\infty} E[\vec{e}_k^T \vec{e}_k] \mu_k^2 < \infty$  almost surely.

**Objective function conditions:** The function  $f(\vec{x})$  is differentiable for each  $x_i$  and is either concave

or unimodal.

**Theorem 1:** For the SP algorithm, if the conditions enumerated above are true, if  $\vec{x}_k$  denotes the solution in the  $k$ th iteration, and if  $\vec{x}^*$  denotes the optimal solution,  $\lim_{k \rightarrow \infty} \vec{x}_k = \vec{x}^*$  with probability 1.

## 2.3 Simulated Annealing

As mentioned above, we have also used a standard meta-heuristic, namely, Simulated Annealing (SA), in our computational experiments. This is a widely-used meta-heuristic that can generate high-quality solutions in problems of combinatorial optimization. It is often claimed that it can escape local optima; see Lundy and Mees (1986) for a convergence analysis. In what follows, we present a quick description of the steps we used for computational experiments.

**Steps.** Start at an arbitrary feasible solution, denoted by  $\vec{x}_{current}$ . Let  $\vec{x}_{best}$  denote the best solution obtained so far, which is initialized to  $\vec{x}_{current}$ . Set  $\psi$ , the so-called temperature, to a high value. The algorithm is run for a number of *phases*. Set  $\bar{P}_{max}$ , the maximum number of phases, to a large number. Each phase, in the algorithm, consists of  $I$  iterations. Within a phase, the temperature is not changed. The values of  $\bar{P}_{max}$  and  $I$  are set according to the time available on the computer. In the algorithm,  $f(\vec{x})$  will denote the objective function value associated with the vector  $\vec{x}$ . Set neighbor search parameter,  $\kappa$ , to a suitable value.

*Step 1.* Set  $\bar{P}$ , the phase number, to 1.

*Step 2.* Randomly select a neighbor of  $\vec{x}_{current}$  as follows.

To select a neighbor, generate a random number  $u(i)$  for  $i = 1, 2, \dots, n$  from the uniform distribution  $U(0, 1)$ . For each  $i = 1, 2, \dots, n$ , do the following: If  $u(i) < 0.5$ ,  $x_{new}(i) \leftarrow x_{current}(i) + \kappa$ ; otherwise  $x_{new}(i) \leftarrow x_{current}(i) - \kappa$ . If any  $x_{new}(i)$  exceeds the upper limit of the feasible value of booking limits or falls below 0, it is projected back into the feasible set  $\mathcal{X}$  (similar to the procedure of SP). In our experiments,  $\kappa$  was set at 3.

*Step 3.* If  $f(\vec{x}_{new}) > f(\vec{x}_{best})$  then set:  $\vec{x}_{best} \leftarrow \vec{x}_{new}$ .

Let  $\Delta = f(\vec{x}_{current}) - f(\vec{x}_{new})$ . If  $\Delta \leq 0$ , set  $\vec{x}_{current} \leftarrow \vec{x}_{new}$ .

Otherwise, that is, if  $\Delta > 0$ , generate a random number  $\bar{u}$  from  $U(0, 1)$ . If  $\bar{u} \leq \exp(-\Delta/\psi)$ , set:  $\vec{x}_{current} \leftarrow \vec{x}_{new}$ .

*Step 4.* Steps 2 and 3 constitute an iteration of the algorithm. Repeat Steps 2 and 3 for  $I$  iterations and then reduce the temperature as follows:  $\psi \leftarrow \bar{G}(\psi)$ , where  $\bar{G}(\psi)$  is a decreasing function of  $\psi$ , e.g.,  $\psi \leftarrow \psi - 10$  or  $\psi \leftarrow \psi/2$ . Then increment  $\bar{P}$  by 1. If  $\bar{P} < \bar{P}_{\max}$ , return to Step 2. Otherwise stop and return  $\vec{x}_{best}$  as the solution.

The effect of simulation noise in evaluating the function *cannot* be ignored. This effect can be analyzed, mathematically, via a result in Gosavi (2002). The result states that by selecting a sufficiently large number of samples it is possible to ensure that the algorithm that uses *noisy* function values mimics the algorithm that uses *exact* function values.

**Theorem 2.** With probability 1, the version of the SA algorithm that uses simulation-based estimates of the function can be made to mimic the version that uses exact function values.

### 3 Computational Results

We begin this section with a description of the EMSR-b heuristic which is used in single-leg and also in network problems. Thereafter, we present computational results with the single-leg and network problem. In the single-leg problems, we use the Poisson distribution for modeling the arrival process. In the network problems, we have used non-homogeneous Poisson processes to model the arrival of customers for booking. We note that our simulation-based approach is independent of the nature of the arrival processes and that these choices were dictated by the need to show that our approach works well with both homogeneous and non-homogeneous Poisson processes. The discrete-event simulator for the single-leg case and the network was coded in C using the approach outlined in Law and Kelton (1999) (see Chapter 2).

#### 3.1 EMSR-b

We note that in our notation:  $V_1 < V_2 < \dots < V_n$ . Let  $Y_i$  denote the random demand for the  $i$ th product in the entire time horizon  $H$ . Then,  $\bar{Y}_i = \sum_{j=i}^n Y_j$  denotes the sum of the demands of all classes above  $i$  and including  $i$ . Then, we define the aggregate revenue for the  $i$ th class to be the weighted mean revenue of all classes above  $i$  and including  $i$  as follows:  $\bar{V}_i = \sum_{j=i}^n V_j E[Y_j] / \sum_{j=i}^n E[Y_j]$ . Then Littlewood's equation of EMSR-b can be given as:  $V_i = \bar{V}_{i+1} Pr(\bar{Y}_{i+1} > S_{i+1})$  for  $i = 1, 2, \dots, n-1$ , where  $S_i$ , the protection level, denotes the number of seats to protect for fare classes  $i, i+1, \dots, n$ . There is no protection level for the lowest class 1. Then, if  $C$  denotes the capacity of the plane, the booking limit for the  $i$ th class is defined to

be:  $BL_i = \max\{C - S_{i+1}, 0\}$  for  $i = 1, 2, \dots, n - 1$ . The booking limit for the highest class  $n$  is the capacity of the plane. If the mean cancellation probability is known to be  $q$ , then cancellations can be accounted for in Littlewood's equation by replacing  $C$  in the above equation by  $C/(1 - q)$  as suggested in Belobaba (1989). Here  $(1 - q)$  is the so-called correction factor. For solving Littlewood's equation, one requires the distribution of each random variable  $\tilde{Y}_i$ .

### 3.2 Single-leg problems

In Table 1,  $Pr_1(\cdot)$  and  $Pr_2(\cdot)$  denote two patterns for arrival probabilities,  $CP_1(\cdot)$  and  $CP_2(\cdot)$  denote two patterns for cancellation probabilities,  $NP(\cdot)$  denotes a pattern for no-show probabilities,  $V(\cdot)$  denotes a pattern used for the revenue per passenger (fare), and  $M(\cdot)$  denotes the cancellation penalty per passenger.

$C_1 = 100$  and  $C_2 = 200$  denote two patterns for the capacity of the plane, and  $\lambda_1(\cdot)$  and  $\lambda_2(\cdot)$  represent two arrival patterns, where  $\lambda_i(l)$  denotes the Poisson rate of arrival in the  $l$ th time zone when the  $i$ th pattern is used. The arrival rate in the  $v$ th class in the  $l$ th time zone is equal to  $Pr_m(v)\lambda_m(l)$ , where  $v \in \{1, 2, \dots, 10\}$ ,  $m \in \{1, 2\}$ , and  $l \in \{1, 2, 3\}$ . Also, the bumping cost is 700, while  $\lambda_1(1) = 0.8$ ,  $\lambda_1(2) = 1.0$ ,  $\lambda_1(3) = 2$ ,  $\lambda_2(1) = 1.5$ ,  $\lambda_2(2) = 2.2$ , and  $\lambda_2(3) = 3.5$ .

Table 1: Data for  $Pr_1(\cdot)$ ,  $Pr_2(\cdot)$ ,  $CP_1(\cdot)$ ,  $CP_2(\cdot)$ ,  $NP(\cdot)$ ,  $V(\cdot)$  and  $M(\cdot)$ ;  $v$  denotes product.

$v$	1	2	3	4	5	6	7	8	9	10
$Pr_1(v)$	0.19	0.17	0.15	0.13	0.11	0.09	0.07	0.05	0.03	0.01
$Pr_2(v)$	0.17	0.15	0.14	0.13	0.10	0.09	0.07	0.06	0.05	0.04
$CP_1(v)$	0.030	0.030	0.035	0.040	0.050	0.070	0.080	0.100	0.150	0.2
$CP_2(v)$	0.050	0.060	0.065	0.080	0.10	0.125	0.15	0.2	0.25	0.2
$NP(v)$	0.001	0.001	0.001	0.001	0.001	0.002	0.009	0.015	0.080	0.1
$V(v)$	100	180	200	240	280	300	330	350	380	400
$M(v)$	70	65	60	55	50	45	40	35	30	25

In Table 3, we present the 95% confidence intervals of the average revenue per flight obtained from simulating the admission policies prescribed by each of the three methods: EMSR-b (van Ryzin and McGill, 2000) (corrected with an overbooking factor), SA, and SP. For EMSR-b calculations, the mean of the arrival rates of the Poisson processes in the three time zones is used as the mean arrival rate of a single Poisson process, and the latter is approximated by the normal distribution to solve Littlewood's equation. We do *not* find any overlap with the results from EMSR-b, in any of the 8 systems studied, and this implies that EMSR-b

Table 2: The patterns used for each system studied.

System	Patterns
1	$Pr_1, CP_1, C_1$ , and $\lambda_1(\cdot)$ .
2	$Pr_2, CP_1, C_1$ , and $\lambda_1(\cdot)$ .
3	$Pr_2, CP_2, C_1$ , and $\lambda_1(\cdot)$ .
4	$Pr_1, CP_2, C_1$ , and $\lambda_1(\cdot)$ .
5	$Pr_1, CP_1, C_2$ , and $\lambda_2(\cdot)$ .
6	$Pr_2, CP_1, C_2$ , and $\lambda_2(\cdot)$ .
7	$Pr_2, CP_2, C_2$ , and $\lambda_2(\cdot)$ .
8	$Pr_1, CP_2, C_2$ , and $\lambda_2(\cdot)$ .

has been outperformed — at least in a statistical sense. In Table 3, the parameter  $IM_{Mt}$  for method  $M$  denotes the percentage improvement of method  $Mt$  over the heuristic (EMSR-b in this case), and is defined as follows:

$$IM_{Mt} = \frac{\bar{R}_{Mt} - \bar{R}_{HEURISTIC}}{\bar{R}_{HEURISTIC}} \times 100, \quad (4)$$

where  $\bar{R}_{Mt}$  denotes the average revenue obtained from applying method  $Mt$  on the problem.

### 3.3 Network problems

For the network problem, we simulated the entire network in the same computer program. Then the simulator was connected to an optimizer — also in the same computer program. The simulator for the entire network, obviously, requires more time for function evaluation in comparison to that written for a single leg. However, we found that optimization can be performed easily, even on small computers available in a university setting, within 15 minutes for a network with 24 legs. With more powerful super-computers, optimization should take an even smaller amount of time. To generate the best results in a given amount of time, we used a *sequential* combination of SP and SA in the optimization process. SP was used first because it quickly produced a good solution. The best solution produced by SP was used as a starting solution for SA. The solution produced by a related linear program (LP), which we will describe below, was used as one of the starting solutions for the SP.

We next describe the DAVN-EMSR-b approach in some detail. This will serve as a benchmark heuristic for our simulation-optimization procedure.

Table 3: Results obtained from EMSR-b, SP, and SA.  $IM$  is defined in (4).  $UL$  and  $LL$  denote the upper and lower confidence interval limits, respectively, on the average revenue in dollars per flight using a 95% confidence level.

System	EMSR-b ( $LL, UL$ )	SP ( $LL, UL$ )	$IM_{SP}$ %	SA ( $LL, UL$ )	$IM_{SA}$ %
1	(22140,22349)	(22595,22899)	2.26	(22546,22806)	1.94
2	(23040,23326)	(23899,24214)	3.77	(23737,23996)	2.95
3	(22768,22959)	(23325,23611)	2.64	(23378,23625)	2.79
4	(21701,21870)	(22209,22409)	2.40	(22158,22560)	2.63
5	(44629,44937)	(45177,45643)	1.40	(45101,45521)	1.18
6	(46816,47142)	(47529,48087)	1.76	(47502,47957)	1.60
7	(45691,46002)	(46537,46896)	1.90	(46561,47031)	2.07
8	(42752,42992)	(44324,44771)	3.91	(44142,44648)	3.55

**DAVN-EMSR-b:** Let  $E[Y_j]$  denote the expected demand and  $V_j$  the revenue for the  $j$ th product. Then the following linear program is solved:

$$\text{Maximize } \sum_{j=1}^n V_j z_j, \text{ such that } 0 \leq z_j \leq E[Y_j], \quad j = 1, 2, \dots, n, \quad (5)$$

$$\text{and } \sum_{j \in \mathcal{A}_l} z_j \leq C^l, \quad l = 1, 2, \dots, L, \quad (6)$$

where  $\mathcal{A}_l$  denotes the set of fare classes that use leg  $l$ , and  $C^l$  denotes the capacity of the plane on the  $l$ th leg. The value of  $z_j$  could be used as booking limit for product  $j$ , and will be used as the starting solution for SA. But a more sophisticated approach that exploits EMSR-b on every leg from the dual prices of the above LP can be employed. This approach is called DAVN-EMSR-b. The *displacement adjusted revenue* (or virtual revenue) for the  $j$ th product that uses leg  $l$ , i.e.,  $DARE_j^l$ , is computed using:

$$DARE_j^l = V_j - \sum_{i \neq l} B_i, \text{ where } j \in \mathcal{A}_i, \quad i \in \{1, 2, \dots, L\} \text{ and } l \in \{1, 2, \dots, L\},$$

and  $B_i$  denotes the dual prices associated with the  $i$ th capacity constraint (6) in the linear program (5). Then  $DARE_j^l$  is treated as the *virtual* revenue of the product  $j$  on leg  $l$ . If there are too many DARE values in a given leg, similar DARE values are clustered (van Ryzin and Vulcano, 2003) to produce a manageable number of aggregate classes. Now that the (virtual) revenue of each product on each leg is available, EMSR-b can be employed on *each leg separately*. For this, on every leg, products that are relevant have to be

re-ordered according to their DARE values; the higher the DARE value, the higher the class. This leads to the generation of separate booking limits for each product-leg combination. The booking limit for product  $j$  on leg  $l$  will be represented as  $BL_j^l$ . A customer requesting a given product is accepted *only if* the conditions with respect to *all* the booking limits are satisfied, i.e., if at time  $t$  in the booking horizon,  $\phi_j(t)$  denotes the number of seats sold for product  $j$ , then a product  $j$  is accepted if  $\phi_j(t) < BL_j^l$  for every leg  $l$  used by product  $j$ . Otherwise that customer is rejected. It could happen that a customer meets the above condition for one leg but not for some other leg; but if the conditions for all legs are not met, the customer is rejected.

We now describe a network of four cities (see Figure 1) that we used for computational purposes. The hub will be denoted by A and the three other cities by X, Y, and Z. Some of the network data are supplied in Tables 4, 5, and 6. Some other data are as follows: the arrival process, which is non-homogeneous Poisson, is described by  $(a = 9, b = 0.03)$  per day, the booking horizon is 100 days long, the penalty of cancellation to the customer is 80 dollars and the penalty to the airline for bumping a passenger is 500 dollars.

For the arrival pattern, we use a non-homogeneous Poisson process, whose intensity function for the time horizon of length  $H$  is defined as  $a + bt$ , where  $t$  denotes the time. We assume concurrent arrivals of all classes and products since this is the most general of assumptions. We must point out that since we use simulation, any arrival distribution can be used just easily. The parameters  $a$  and  $b$  for the  $i$ th product equal  $Pr(i)a$  and  $Pr(i)b$ , respectively, such that each product has its own independent non-homogeneous Poisson process (Ross, 2003). For simulating the non-homogeneous process, we used the method described as “Method 1” on page 59 of Kao (1997).

For EMSR-b calculations, the integrated intensity function of the non-homogeneous Poisson process,  $m(t)$ , was calculated so that one could use a Poisson approximation (see page 57 of Kao (1997)); the Poisson process could be further approximated in a convenient fashion by the normal distribution to solve Littlewood’s equation. For our function, the integrated intensity function turns out to be  $m(t) = at + 0.5bt^2$  from which the mean demand in the time horizon of length  $H$  becomes  $aH + 0.5bH^2$ . This is used as the mean (and also the variance) of an equivalent Poisson distribution, which, as stated above, can be approximated easily (and accurately) by the normal distribution. For the simulation-optimization approach, the products are ranked using DAVN values, i.e., by their net worth, which was calculated as  $W_j = \sum_{l \in \mathcal{L}_j} DAVN_j^l$ , where  $\mathcal{L}_j$  denotes the set of legs needed for product  $j$ . The higher the value of  $W$  for a product, the higher its rank.

For the SP-SA combination, the booking limits derived from the LP are used as a starting solution. Fare structures defined in Table 6 are used to define the five networks studied here. Table 7 shows the actual revenues obtained from DAVN-EMSR-b and the SP-SA combination. As is clear, the SP-SA combination is

clearly superior to DAVN-EMSR-b. We must add that in some cases, SA did not improve at all upon the solution of SP, and hence the solution from the combination is essentially that of SP. In our experiments with SP, we used  $w = 0.5$  and  $\mu^{small} = 0.01$ .

Table 4: Description of various legs in the network.

Plane (leg)	Origin	Destination	Capacity
1	A	X	100
2	X	A	100
3	A	Y	100
4	Y	A	100
5	A	Z	100
6	Z	A	100

## 4 Conclusions

This paper presented an integrated simulation-based approach that can be applied to solve a complex seat-allocation problem in the airline industry. The model developed accommodated most real-life considerations, including cancellations and overbooking, which are ignored in many models in the literature. Two efficient optimization techniques were combined with the simulation model. Computational results showed that our simulation-optimization approach can outperform both EMSR-b for single-leg problems and DAVN-EMSR-b for network problems. The single-leg results are from the MS thesis of Ozkaya (2002). A further improvement over this approach could be realized by simulating each leg on a separate processor in a parallel-processing environment. Future research will be directed towards minimizing the run time by using parallelization techniques.

**Acknowledgements:** This work was partly supported from grant DMI: 0114007 to the first author from the National Science Foundation. The authors express gratitude to the two anonymous reviewers and the special-issue guest editor, Prof. Alf Kimms, for their useful suggestions.

Table 5: The meanings of the symbols are as follows.  $Pr(i)$ : Probability that an arriving passenger requests the  $i$ th product, and  $CP(i)$ : Probability of cancellation for the  $i$ th product.

Product $(i, j)$	Itinerary	$(Pr(i), Pr(j))$	$(CP(i), CP(j))$
(1,13)	A $\rightarrow$ X	(0.056,0.014)	(0.025,0.3)
(2,14)	X $\rightarrow$ A	(0.064,0.016)	(0.025,0.3)
(3,15)	A $\rightarrow$ Y	(0.048,0.012)	(0.025,0.3)
(4,16)	Y $\rightarrow$ A	(0.056,0.014)	(0.05,0.3)
(5,17)	A $\rightarrow$ Z	(0.064,0.016)	(0.05,0.3)
(6,18)	Z $\rightarrow$ A	(0.048,0.012)	(0.075,0.3)
(7,19)	X $\rightarrow$ Y via A	(0.08,0.020)	(0.125,0.3)
(8,20)	Y $\rightarrow$ X via A	(0.096,0.024)	(0.2,0.3)
(9,21)	X $\rightarrow$ Z via A	(0.08,0.020)	(0.2,0.3)
(10,22)	Z $\rightarrow$ X via A	(0.072,0.018)	(0.225,0.3)
(11,23)	Y $\rightarrow$ Z via A	(0.08,0.02)	(0.2,0.3)
(12,24)	Z $\rightarrow$ Y via A	(0.056,0.014)	(0.2,0.3)

Table 6: The table enumerates a number of fare structures for the network problem.  $FS_k$  denotes the  $k$ th fare structure.

Product	$FS_1$	$FS_2$	$FS_3$	$FS_4$	$FS_5$
$(i, j)$	$(V_i, V_j)$	$(V_i, V_j)$	$(V_i, V_j)$	$(V_i, V_j)$	$(V_i, V_j)$
(1,13)	(350,700)	(250,500)	(350,500)	(250,400)	(125,250)
(2,14)	(375,750)	(275,550)	(375,525)	(275,425)	(175,350)
(3,15)	(400,800)	(300,600)	(400,550)	(300,450)	(200,400)
(4,16)	(430,860)	(330,660)	(430,585)	(330,480)	(230,460)
(5,17)	(450,900)	(350,700)	(450,600)	(350,500)	(250,500)
(6,18)	(500,1000)	(400,800)	(500,650)	(400,550)	(300,600)
(7,19)	(600,1200)	(500,1000)	(600,750)	(500,650)	(350,700)
(8,20)	(610,1220)	(510,1020)	(610,760)	(510,660)	(375,750)
(9,21)	(620,1240)	(520,1040)	(620,770)	(520,670)	(380,760)
(10,22)	(630,1260)	(530,1060)	(630,780)	(530,680)	(390,780)
(11,23)	(640,1280)	(540,1080)	(640,790)	(540,690)	(395,790)
(12,24)	(650,1300)	(550,1100)	(650,800)	(550,700)	(400,800)

Table 7:  $\bar{R}_{Mt}$  denotes the expected revenue in dollars per flight when method  $Mt$  is used.  $LP^*$  denotes the solution of (5), which forms an upper bound on the network revenues.

Fare Structure	$\bar{R}_{DAVN-EMSR-b}$	$\bar{R}_{SP-SA}$	$IM_{SP-SA}$	$LP^*$
1	283739.03	299723.34	5.63	337136
2	227289.78	258984.27	13.94	268803
3	203190.75	229602.85	12.99	258269
4	168539.17	190217.56	12.86	208596
5	144905.63	156285.29	7.85	188706.5

## References

- P.P. Belobaba. Application of a probabilistic decision model to airline seat inventory control. *Operations Research*, 37:183–197, 1989.
- D.P. Bertsekas and J.N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.
- D. Bertsimas and S. de Boer. Simulation-based booking limits for airline revenue management. *Operations Research*, 53(1):90–106, 2005.
- D. Bertsimas and I. Popescu. Revenue management in a dynamic network environment. *Transportation Science*, 37(3):257–277, 2003.
- J.F. Bonnans and A. Shapiro. *Perturbation Analysis of Optimization Problems*. Springer, New York, 2000.
- E.A. Boyd and I. Bilegan. Revenue management and e-commerce. *Management Science*, 49(10):1363–1386, 2003.
- S.L. Brumelle and J.I. McGill. Airline seat allocation with multiple nested fare classes. *Operations Research*, 41:127–137, 1993.
- R. E. Chatwin. Multiperiod airline overbooking with a single fare class. *Operations Research*, 46 (6):805–819, 1998.
- V. Chen, D. Gunther, and E. Johnson. Solving for an optimal airline yield management policy via statistical learning. *Journal of the Royal Research Society: Series C*, 52(1):19–30, 2003.
- W. Cooper and T. Homem-de-Mello. Revenue management using sampling-based optimization and Markov decision processes. Working paper at University of Minnesota and Northwestern University, 2004.
- R.E. Curry. Optimal airline seat allocation with fare classes nested by origins and destinations. *Transportation Science*, 24:193–204, 1990.
- S. de Boer, R. Freling, and N. Piersma. Stochastic programming for multiple-leg network revenue management. Working paper with ORTEC Consultants, Gouda, Netherlands, 2004.
- M. Fu and S. Hill. Optimization of discrete-event systems via simultaneous perturbation stochastic approximation. *IEEE Transactions*, 29(3):233–243, 1997.
- F. Glover. Tabu search: A tutorial. *Interfaces*, 20:74–94, 1990.

- F. Glover, R. Glover, J. Lorenzo, and C. McMillan. The passenger mix problem in scheduled airlines. *Interfaces*, 12:73–79, 1982.
- A. Gosavi. The effect of noise on artificial intelligence and meta-heuristic techniques. In *Proceedings of the Artificial Neural Networks in Engineering Conference (Intelligent Engineering Systems Through Artificial Neural Networks)*, volume 12, pages 981–988. American Society of Mechanical Engineering Press, 2002.
- A. Gosavi. *Simulation-Based Optimization: Parametric Optimization Techniques and Reinforcement Learning*. Kluwer Academic Publishers, Boston, MA, 2003.
- A. Gosavi. A reinforcement learning algorithm based on policy iteration for average reward: Empirical results with yield management and convergence analysis. *Machine Learning*, 55(1):5–29, 2004.
- A. Gosavi, N. Bandla, and T. K. Das. A reinforcement learning approach to a single leg airline revenue management problem with multiple fare classes and overbooking. *IIE Transactions*, 34:729–752, 2002.
- J. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, MI, 1975.
- R. Howard. *Dynamic Probabilistic Systems Volume II: Semi-Markov Decision Processes*. John Wiley and Sons, NY, USA, 1971.
- E. Kao. *An Introduction to Stochastic Processes*. Duxbury Press, Belmont, CA, USA, 1997.
- I. Karaesmen and G. van Ryzin. Overbooking with substitutable inventory classes. *Operations Research*, 52(1):83–104, 2004.
- S. Kirkpatrick, C. Gelatt, and M. Vecchi. Optimization by simulated annealing. *Science*, 220:671–680, 1983.
- A. M. Law and W. D. Kelton. *Simulation Modeling and Analysis, Third Edition*. McGraw Hill, Inc., New York, NY, 1999.
- T.C. Lee and M. Hersh. A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science*, 27:252–265, 1993.
- K. Littlewood. Forecasting and control of passenger bookings. In *Proceedings of the 12th AGIFORS (Airline Group of the International Federation of Operational Research Societies) Symposium*, pages 95–117, 1972.
- M. Lundy and A. Mees. Convergence of the annealing algorithm. *Mathematical Programming*, 34:111–124, 1986.

- J.I. McGill and G. J. van Ryzin. Revenue management: Research overview and prospects. *Transportation Science*, 33(2):233–256, 1999.
- E. Ozkaya. A simulation-based approach to seat allocation in airplanes. Unpublished MS thesis at Colorado State University-Pueblo, Colorado, USA, 2002.
- L. W. Robinson. Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes. *Operations Research*, 43:252–263, 1995.
- S. Ross. *Introduction to Probability Models*. John Wiley and Sons, New York, NY, eighth edition, 2003.
- E. Shlifer and Y. Vardi. An airline overbooking policy. *Transportation Science*, 9:101–114, 1975.
- R.W. Simpson. Using network flow techniques for market and seat inventory control. MIT Flight Transportation Laboratory, Cambridge, MA, 1989.
- B.C. Smith and C.W. Penn. Analysis of alternative origin-destination control strategies. In *Proceedings of the AGIFORS Symposium (Sydney, Australia)*. 1988.
- J. Spall. Multivariate stochastic approximation using a simultaneous perturbation gradient approximation. *IEEE Transactions on Automatic Control*, 37(3):332–341, 1992.
- J. C. Spall. *Introduction to Stochastic Search and Optimization*. John Wiley and Sons, NJ, USA, 2003.
- J. Subramaniam, S. Stidham Jr, and C.J. Lautenbacher. Airline yield management with overbooking, cancellations and no-shows. *Transportation Science*, 33(2):147–167, 1999.
- K. Talluri and G. van Ryzin. *The Theory and Practice of Revenue Management*. Kluwer Academic, Boston, MA, 2004.
- G. van Ryzin and J. I. McGill. Revenue management without forecasting or optimization: An adaptive algorithm for determining seat protection levels. *Management Science*, 46:568–573, 2000.
- G. van Ryzin and G. Vulcano. Simulation-based optimization of virtual nesting controls of network revenue management. Working paper series at Columbia University, NY, 2003.
- B. Vinod. Origin-and-destination yield management. In *The Handbook of Airline Economics (edited by D. Jenkins)*, pages 459–468. The Aviation Weekly Group of the McGraw-Hill Company, New York, 1995.
- E.L. Williamson. Airline network seat inventory control: Methodologies and revenue impacts. Unpublished Ph.D. Dissertation Flight Transportation Lab, MIT, Cambridge, 1992.

- R.D. Wollmer. An airline seat management model for a single leg route when lower fare classes book first. *Operations Research*, 40:26–37, 1992.
- J.T. Wong, F.S. Koppelman, and M.S. Daskin. Flexible assignment approach to itinerary seat allocation. *Transportation Science*, 27B:33–48, 1993.