Operations Research

In This Issue

OPERATIONS RESEARCH
Vol. 53, No. 1, January–February 2005, pp. ii–iv
ISSN 0030-364X | EISSN 1526-5463 | 05 | 5301 | 00ii

DOI 10.1287/opre.1040.0177

In This Issue

Measuring the Achievement of Objectives

In all aspects of our field, operations researchers use models of all kinds, including math programming, inventory control, supply chain management, queuing, public systems, cost-benefit analysis, and decision analysis. We use these models to design and evaluate alternative courses of action. This requires, among other things, measures that can be used to describe the consequences of the various alternatives in terms of stated objectives. Too often, it seems as if the chosen measures do not provide the basis for an adequate description. As a result, the models cannot provide the hoped-for insight. In "Selecting Attributes to Measure the Achievement of Objectives," R. L. Keeney and R. S. Gregory lay out a systematic procedure based on commonsense logic to select a useful set of measures, referred to as attributes, for the objectives stated for any model. The paper covers both the selection of existing attributes and the design of constructed attributes for objectives with no appropriate existing measure.

Dynamic Scheduling in Complex Communication Systems

Many elements of modern communication networks are very different from a simple "data link" of fixed capacity. As an example, a wireless channel shared by multiple traffic flows (users) is often such that feasible instantaneous transmission rates of different users are (a) mutually interdependent and (b) randomly varying in time. The problem of dynamic scheduling of users' data transmissions in such complex systems to optimize certain performance objectives has received much attention in recent years. In "On the Asymptotic Optimality of the Gradient Scheduling Algorithm for Multiuser Throughput Allocation," A. L. Stolyar considers the problem of maximizing system aggregate "utility," which is a concave function of users' average throughputs. The paper rigorously proves that a rather "myopic" gradient-type dynamic scheduling algorithm is in fact (asymptotically) optimal for this problem. Moreover, the analysis of the user throughput dynamics demonstrates that the algorithm is very robust in that it can quickly adjust to changes in system parameters. The gradient algorithm is very attractive for applications: a special case of it is employed in some commercial cellular data networks.

Launching New Products in the Presence of Competition

Companies launching new products or services are often forced to maintain a balance between a set of conflicting objectives, including time-to-market, product unit cost, and development cost. To find a balance between these conflicting objectives, teams typically rely on simple trade-off rules. While understanding the financial impact of changes in development cost and product unit cost is relatively simple, understanding the financial impact of a change in launch time, which affects, among other things, the competitive dynamics for the product, has proven to be rather difficult. In "Optimal Product Launch Times in a Duopoly: Balancing Life-Cycle Revenues with Product Cost," S. Savin and C. Terwiesch propose a model that allows a quantification of the impact of the time-to-market decision on the overall profitability of a new product in the presence of competition. Their research develops a set of practically important quantitative guidelines for the analysis of time-to-market decisions. For example, the authors show that the simplistic marginal cost-benefit analysis often used in practice may lead to significant losses in lifecycle profits.

Stochastic Supply Chain Design Problem

In today's competitive market, a company's distribution network must meet service goals at the lowest possible cost. Relevant costs include inbound and outbound transportation costs, fixed and variable warehouse costs, inventory costs from carrying and producing or sourcing from different locations. Complex trade-offs in the distribution network system make these costs difficult to analyze. The distribution network structure and warehouse-retailer assignments have immediate impact on the replenishment policies and costs at the warehouses and retail outlets. To understand the trade-offs, we need to incorporate models from stochastic inventory theory into the network design domain. In "Stochastic Transportation-Inventory Network Design Problem," J. Shu, C.-P. Teo, and Z.-J. M. Shen address these issues and propose a practical and efficient column generation algorithm to solve this stochastic supply chain design problem.

Continuous Location Problems

What is the best place for a facility with respect to a given demand distribution? Such problems of finding an optimal location for a store, a factory, or some other facility are often crucial for the success of a business. The goal is to find a point that minimizes some measure of distance to a set of demand points. In the classical Fermat-Weber or median problem, this measure is the average distance to a finite set of points. The continuous Fermat-Weber problem arises when the demand is modeled as a continuous set with an associated probability distribution. Additional difficulties arise from the presence of "obstacles," which affect the distances between points. In "On the Continuous Fermat-Weber Problem," S. P. Fekete, J. S. B. Mitchell, and K. Beurer give a first exact algorithmic study of the continuous Fermat-Weber problem for the case in which distances are measured according to the Manhattan metric, possibly among obstacles. Among their results are various polynomial-time algorithms for different versions of the problem, as well as a proof that the problem gets substantially more difficult when many facilities have to be placed. Their study makes strong use of methods from computational geometry.

Managing Navy Personnel

There are about 300,000 sailors in the United Sates Navy, and managing their careers is challenging. In good economic times, many well-trained sailors leave the navy for higher paying jobs in corporate America. In an attempt to retain experienced sailors, the navy is investigating how to increase a sailor's opportunities within the navy. The idea is to allow a sailor to choose their next job from a preselected list of positions for which the sailor is qualified. This differs from the current practice of assigning sailors to jobs that are deemed most important by the navy. The Hearin Center for Enterprise Science at the University of Mississippi was charged with finding a technique to generate "optimal" job lists. Besides ensuring naval policy, an optimization model needed to balance the personal needs of the navy with the individual goals of a sailor. In "Navy Personnel Planning and the Optimal Partition," A. Holder worked with the sailors in charge of job reassignments at Millington Naval Base and designed an optimization model that highlights desirable assignments. This optimization model also measures competition to further aid the creation of a job list.

Simulation as a Powerful Tool for Airline Revenue Management Optimization

Airlines offer a large number of itineraries, both direct and connecting, each at widely different fares depending on the ticket's terms and conditions. All these "products" share the same resources, seats on flights, so the airline has to decide how many of each of these to make available to maximize expected revenue based on predicted demand. This decision process is known as revenue management and has a large impact on the airline's bottom line. Mathematical optimization techniques have allowed the airlines to develop good policies, but because of the complexity of the problem they have not been very sophisticated and are typically by necessity based on unrealistic simplifying assumptions. In "Simulation-Based Booking Limits for Airline Revenue Management," D. Bertsimas and S. de Boer develop a simulation-based optimization method based on a more realistic representation of the airline booking process. The revenue improvements they show are possible will hopefully motivate further research into the use of simulation-based optimization in this field, which is a powerful and promising technique because of its modeling flexibility.

Stability of Data Networks

Is the bandwidth resource in the Internet properly utilized? Which bandwidth sharing policies will lead to satisfactory network performance in terms of, for example, transmission throughout? In designing control protocols (e.g., the Transmission Control Protocol (TCP)) for the Internet, it is important to gain better understanding of problems such as those mentioned above. In "Stability of Data Networks: Stationary and Bursty Models," H.-Q. Ye, J. Ou, and X.-M. Yuan apply the fluid network model approach to study the stability of the connection level data network that captures some macroscopic features of data communication networks including the Internet. They show that, under many bandwidth sharing policies discussed widely in the literature, the network model is stable if the average offered transmission load on each transmission link is within its bandwidth capacity for data transmission. On the other hand, they also find an interesting counterintuitive example for a bandwidth sharing policy related to a widely used TCP.

Efficient Simulation-Based Solution of Stochastic Dynamic Programs

The analysis of many modern systems necessitates stochastic simulation, for which dynamic programming can provide a framework for sequential decision making under uncertainty. For example, in a capacity planning model in manufacturing, the transitions and cost/rewards in a dynamic programming model might be estimated from output runs of a large simulation model of a semiconductor fabrication facility, and the action might be a choice of whether to add long-term capacity by purchasing an expensive new piece of machinery. However, simulation can be computationally expensive, and numerical solution of the resultant dynamic program can become prohibitive for large state spaces. In "An Adaptive Sampling Algorithm for Solving Markov Decision Processes," H. S. Chang, M. C. Fu, J. Hu, and S. I. Marcus propose an algorithm for finding the optimal value of a finite horizon dynamic program where only samples of the stochastic behavior are available. Adapting theory from multiarmed bandit problems, the algorithm adaptively chooses which action to sample next, based on the results obtained from samples taken up to that point. This adaptive approach has the potential to yield substantial computational efficiency savings in settings where the sampling cost is relatively expensive, since it avoids wasting computational resources on sampling prospectively poor actions.

Two-Staged Two-Dimensional Cutting: An Exact Algorithm

In many industrial sectors, in particular in the packing industry, minimizing waste is a critical problem that is frequently encountered. In "An Exact Algorithm for Constrained Two-Dimensional Two-Staged Cutting Problems," M. Hifi and R. M'Halla propose an exact algorithm for the constrained two-dimensional packing problem where a set of rectangular pieces are to be cut on a rectangular stock sheet of fixed width and length. The objective is to identify a cutting pattern that maximizes the usage of the rectangular stock sheet while respecting the upper demand value for each piece type. Each piece is produced in the final cutting pattern by at most two cuts that are perpendicular to the edges of the stock sheet.

The Science of Variety

Virtually all manufacturers and distributors produce and distribute multiple products (items). Each customer order typically requests several items. Customer satisfaction is order based, relying on the fulfillment of all items in the entire order. Production and procurement planning systems are usually item based, however, focusing on each item separately. The same situation arises in assemble-to-order manufacturing. Filling a demand requires all the components in a product, but procurement metrics focus on one component at a time. Is there a way to improve order-based performance by utilizing the existing item-based systems? How much effort does it require, and how much benefit can be achieved? In "Order-Based Cost Optimization in Assemble-to-Order Systems," Y. Lu and J.-S. Song explore these issues and provide promising results.

Stability Regions in Single-Product Lot-Sizing Problems

Constructing stability regions for the solution of any mathematical problem is a very interesting issue, since the stability region provides the space(s) for the parameters involved in the problem where the obtained optimal solution remains unchanged or it changes in a known way. In practice, one almost never knows the exact values of problem parameters and must resort to their estimates. So the need for an effective tool capable of assessing scenarios that correspond to various parameter sets is obvious. In "The Single-Product Lot-Sizing Problem with Constant Parameters and Backlogging: Exact Results, a New Solution, and All Parameter Stability Regions," I. Ganas and S. Papachristos propose a new algorithm determining stability regions for the single-product lot-sizing problem with constant parameters and backlogging. While the construction of a stability region(s) is a hard mathematical problem, the proposed algorithm is easy to implement and is therefore suitable for managerial applications.