Operations Research

In This Issue

To cite this article:
(2007) In This Issue. Operations Research 55(4):ii-iv. https://doi.org/10.1287/opre.1070.0465

Operations Research
Vol. 55, No. 4, July–August 2007, pp. ii–iv
issn 0030-364X | eissn 1526-5463 | 07 | 5504 | 00ii

doi 10.1287/opre.1070.0465

In This Issue

Profiting Despite Proliferation

As product portfolios continue to grow, through additions of options and features to meet the increasing diverse needs of a global marketplace, Deere applies operations research techniques to optimally design its product line. In "Building Efficient Product Portfolios at John Deere and Company," T. H. Yunes, D. Napolitano, A. Scheller-Wolf, and S. Tayur show how accounting for substitution across products, and even defections to competitors, Deere reaps tens of millions in additional profit annually using a novel customer migration model that drives a computational method using a mixed-integer programming model.

Impact of Outsourcing Component Capacities

In recent years, managers and scholars have explored a variety of methods for leveraging operational flexibility to improve supply chain responsiveness while avoiding excessive inventories. Some examples include the use of common components across multiple finished products, operational hedging based on flexible production capacity, and postponing the point at which a common product is specialized. During this same time period, manufacturers have increasingly outsourced parts of their operations, leading to more decentralized decision structures – for example, by allowing suppliers to manage production capacities or inventories of product components. In "Incentives and Commonality in a Decentralized Multiproduct Assembly System," F. Bernstein, G. A. DeCroix, and Y. Wang analyze the impact of this decentralization trend on aspects of system behavior related to operational flexibility. Decentralizing the decision structure can significantly reduce the use of operational flexibility and lead to several types of inefficiency. For example, it can reduce the breadth of product offerings, reduce the likelihood that component capacities will exploit risk pooling opportunities, and make the use of common components less attractive.

Bid Prices with Higher Fidelity

For many years now, bid-price control has been one of the most popular heuristics for network revenue management problems, such as those faced by airlines. At the same time, new methods have been emerging in approximate dynamic programming (ADP). In "Dynamic Bid Prices in Revenue Management," D. Adelman brings the latest ADP methods to bear on the network revenue management problem. The approach grounds standard bid-price control in dynamic programming methodology and provides the capability to compute more powerful bid prices that depend on time. It also leads to a general, new class of linear programs with arbitrarily higher fidelity.

Building Capacitated Networks to Meet Uncertain Demand

Designing cost-effective networks that can respond to high variability in demand is becoming a major topic of interest in many fields, such as logistics/supply chain, telecommunication, and transportation. Situations in which decisions must be made with limited demand information is especially challenging as one would like to avoid building too much excess capacity, but still meet demand with high probability. In "Two-Stage Robust Network Flow and Design Under Demand Uncertainty," A. Atamtürk and M. Zhang address this problem with a two-stage robust network optimization approach in which one defers a subset of the flow decisions until after the realization of the uncertain demand. The authors discuss applications of their general approach to production and facility location problems and provide computational experience.

Inventory Systems with Fixed Replenishment Intervals

In multiechelon production/inventory supply chains there are batching decisions (or rules) and safety-stock decisions. Batching rules are typically high-level decisions, determined by systemic fixed batch sizes or fixed replenishment intervals. Given these rules, inventory or safety stock decisions are then made, specifying the replenishment strategy used to operationalize the batching rules. In "Optimal Control of Serial Inventory Systems with Fixed Replenishment Intervals," G.-J. van Houtum, A. Scheller-Wolf, and J. Yi study a serial system with a given fixed replenishment interval per stage as the batching rule. These replenishment intervals decrease as we move downstream through the supply chain, as do the setup times and costs. The authors show that base stock policies remain optimal in this setting, and they describe efficient solution procedures to find optimal base stock parameters. This research brings time-based batching rules into supply chain planning solutions, incorporating multiechelon inventory theory.

Looking to the Future When Making a Decision in the Present: How Far Is Far Enough?

The modern manager routinely struggles with the issues associated with acquiring relevant information in order to make an effective decision. The temptation is to seek a great deal of information in order to feel confident that the decision will be robust. It would certainly be valuable to know when the available information is sufficient to make the best possible current decision, in which case seeking additional information would be of no benefit toward improving the initial decision. The concept of a "forecast horizon" was defined more than five decades ago to highlight this "sufficiency" of information in decision-making situations, but its discussion was mostly confined to academic circles. In "Forecast Horizons for a Class of Dynamic Lot-Size Problems Under Discrete Future Demand," M. Dawande, S. Gavirneni, S. Naranpanawe, and S. Sethi present structural and computational results for forecast horizons obtained by assuming that future demand is discrete.

Performance Prediction and Preselection for Optimization and Heuristic Solution Procedures

One of the significant contributions of operations research is the solution of complex planning problems, using optimal or heuristic methods. This research typically involves the justification of new solution procedures by computational testing on large random data sets. By contrast, in "Performance Prediction and Preselection for Optimization and Heuristic Solution Procedures," N. G. Hall and M. E. Posner propose a methodology for using existing solution procedures more effectively. They identify specific characteristics of problem data that significantly affect the performance of a solution procedure. The methodology predicts, with a high degree of probability, which solution procedure is more efficient. This provides guidelines for the choice of an appropriate solution procedure for a given data set. This methodology is applied to a well-known optimization problem for which alternative solution procedures are available. These guidelines result in significant improvements in computational efficiency, without the need to develop new solution procedures.

Leveraging Computing Grids for Decentralized Optimization

The availability of computing grids, whether academic (e.g., Virginia Tech's System X) or commercial (e.g., Sun Microsystems' Sun Grid, Amazon's EC2, etc.), is on the rise. A natural use of such computational power is to solve complex optimization problems. "A Decentralized Approach to Discrete Optimization via Simulation: Application to Network Flow," by A. Garcia, S. D. Patek, and K. Sinha analyzes an approach where computers in the grid take on the role of "players" and engage in an artificial game in order to solve a complex optimization problem. The interaction of these players leads to a decentralized exploration of the search space associated with the optimization problem and the identification of locally optimal solutions. The scheme proposed in the paper lends itself well for cases in which performance measures (or objective function values) are only available from a noisy simulation of the system to be optimized. Hence, the scheme is one in which the search and simulation tasks are decentralized. These ideas are illustrated with an application to the optimization of intradomain network flow, where the scheme proposed is proven to identify the global optimal solution.

Improving Truck Scheduling and Inventory Management

For a German logistics company operating different warehouses that are connected by a truck shuttle, the problem of planning the schedules and truckloads is essential to fulfill all deliveries to its customers in time. On the one hand the cost of the schedule should be minimized; on the other hand the company needs to effectively manage inventory because it is impossible to transport all articles for all deliveries just in time. In "A Case Study of Joint Online Truck Scheduling and Inventory Management for Multiple Warehouses," C. Helmberg and S. Röhl develop a multicommodity network flow model over time and an appropriate piecewise linear convex objective to tackle this problem. The objective is based on estimations of the distributions of the demand for the articles at the different warehouses. The paper contains various real-world examples that show that their approach works in reasonable time, reduces the number of truck rides, and indeed improves inventory management.

Exploiting Flexibility to Compensate for Failures

Over the past few years, there has been significant interest in how operational flexibility may be exploited to improve performance in service and production systems. Most of this work has concentrated on the case when all components are 100% reliable. In "Compensating for Failures with Flexible Servers," S. Andradóttir, H. Ayhan, and D. G. Down study how flexibility can be used to counteract the effects of failures of both workers and the resources they employ. The setting is that of a general queuing network, for which the capacity is found by solving a linear programming problem. The solution of the linear program in turn allows one to easily construct policies that guarantee a given capacity. Furthermore, the authors identify conditions under which flexibility can completely compensate for failures, allowing managers to choose flexibility structures that work well in failure-prone systems.

Calibrating Military Targets

A weapons platform shooting at encountered enemy targets can be thought of as a system offering a service to unwilling recipients. However, the nature of military engagement means that it is a service system with some very unusual characteristics. The shooter/server, which might, for example, be an air-defense missile launcher, is vulnerable to being disabled. Moreover, the extent and speed of any such disabling may well depend upon the transitory task (target) currently assumed. Further, much information about these targets (tasks), such as their likely (or actual) defensive lethality, may be revealed only while "service" is actually in progress. The shooter/server seeks to prioritize targets in accordance with their potential offensive lethality while herself remaining survivable. As part of her strategy she may also choose to temporarily disengage from service if surviving to fight another day demands it. The computational burden of determining good "service" policies is much reduced if it can be shown that there is some way of calibrating the targets currently within range (using all information that is available to the shooter about each as the conflict proceeds) that can support effective decision making. In "Index Policies for Shooting Problems," K. D. Glazebrook, C. Kirkbride, H. M. Mitchell, D. P. Gaver, and P. A. Jacobs discuss when this is the case and how such calibrations are developed.

Where Is the Best Circle?

In such diverse applications as search-and-rescue missions and biological studies on the growth of rats, one may encounter the problem of locating a spherical circle on a sphere. In "Locating a Circle on a Sphere," J. Brimberg, H. Juel, and A. Schöbel address various models of this problem. In the minisum version, the best circle is the one minimizing the sum of the geodesic distances between a number of given points on the surface of the sphere and the circle. In the minimax version, they seek the circle minimizing the maximum of these distances. They provide polynomial algorithms for the minimax model and for the case of locating a great circle in the minisum model. For the case of locating an arbitrary minisum circle, they outline a branch and bound procedure.

Chance Constraints: Basic Theory Too Simple, but Burdens of Complexity May Be Avoided

Chance-constrained programming may or may not involve joint reliability requirements among the constraints, depending on the application. For example, consider a groundwater containment design problem characterized as minimizing the cost of several injection and extraction wells subject to several chance constraints requiring that the hydraulic gradient be inward-directed at locations around the capture zone. If one location fails, capture fails, and it is as if all locations fail. Thus the analyst wishes to impose the reliability requirement embodied in the chance constraints jointly, rather than independently. Such a requirement, however, implies a significant computational burden, as well as analytical challenges. In "A Screening Technique for Joint Chance-Constrained Programming for Air-Quality Management," H. An and J. W. Eheart propose that instead of such full joint analysis, chance constraints may be imposed independently, but with appropriately chosen pessimistic (i.e., higher) reliability requirements. The article offers a screening technique that may be used to determine a first-cut estimate, based on the limits of values of correlations coefficients, of the range of such pessimistic independent reliabilities. Using these endpoints-of-spectrum values, the analyst may determine whether the analytical and computational effort of a full joint analysis is justified.