



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

In This Issue

To cite this article:

(2006) In This Issue. Operations Research 54(1):ii-iv. <https://doi.org/10.1287/opre.1060.0275>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2006, INFORMS

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

In This Issue

Using Expensive Hospital Imaging Equipment More Efficiently

As health care costs continue to grow and patient satisfaction becomes a more critical factor in an increasingly competitive environment, hospitals, which account for over 30% of all health care costs, are under increasing pressure to increase efficiency while improving quality. These tensions are particularly apparent in the management of diagnostic imaging equipment such as MRIs. Hospital managers are under pressure to keep these very expensive machines highly utilized, yet in many hospitals lack of timely access to an MRI leads to backups in the emergency room, disruptions in the operating room schedule, and delays in the discharge of inpatients. Outpatients also experience delays, which can lead to loss of business for the hospital. In “Managing Patient Service in a Medical Diagnostic Facility,” L. V. Green, S. Savin, and B. Wang address the problem of the effective allocation of expensive diagnostic imaging capacity among these competing groups of patients. Based on data from an MRI facility in a large urban teaching hospital, the authors analyze the structure of optimal real-time capacity management policies and identify some simple heuristics that work well for a large range of realistic problem parameters.

Determining How the Possibility of Learning Impacts Decisions

One thing that we know in this uncertain world is that uncertainty and learning have implications for current decisions. This fact is particularly relevant to the debate over what to do about global warming, where the costs of immediate action are high, but the possibility of catastrophic damage looms if society does nothing. In “Increasing Risk and Increasing Informativeness: Equivalence Theorems,” E. Baker provides results that allow one to determine the impact of learning on optimal decisions and on equilibria of games through exploiting the parallel between risk and learning. This paper was inspired by the complexities of the climate change problem, but can be applied wherever uncertainty and learning are important.

Approximations for Multiserver Queues with Abandonments

Customer abandonment is a significant feature of many service systems, such as telephone call centers. In such systems, it is usually appropriate to model both the times

to abandon (before starting service) and the service times as random variables. However, data often reveal that these times are not nearly exponentially distributed, making a model that incorporates the true distributions very difficult to analyze. In “Fluid Models for Multiserver Queues with Abandonments,” W. Whitt develops fluid models that serve as useful approximations, especially for heavily loaded systems.

Decision Aids for Scheduling and Hedging in Electricity Markets

Throughout the world, electricity markets are undergoing dramatic changes. Today, it is common to trade electricity through a variety of transactions such as electricity forwards, fuel (gas) forwards, spotmarket trading, and the like. Due to the volatility of electricity and fuel prices, and the inherent uncertainty associated with electricity load forecasts, producers of electricity are under increasing pressure to balance risk and return. In the paper “A Stochastic Programming Approach to Power Portfolio Optimization,” S. Sen, L. Yu, and T. Genc present a model that provides recommendations for electricity forwards contracts, while accommodating daily fluctuations in electricity loads, spot-market prices, and ensuring that the resulting portfolio avoids positions that expose the producer to the risk of very large losses in the future. The need to incorporate both forward contracts as well as spot-market activity leads to a multiscale stochastic optimization model in which certain aspects of the model capture operational day-to-day realities, whereas other aspects of the model are governed by medium-term considerations. This multiscale approach calls for a new algorithmic method, which makes it possible to address instances resulting from realistic data. The authors provide extensive simulation results to support their claim that their model provides significant advantages over a standard approach.

Integrated Logistics: Simultaneously Installing Facilities and Transportation Networks

When considering where to locate manufacturing facilities or warehouses, firms must also take into account the cost of transporting the goods from the facilities to their final destinations. If there are multiple final destinations and the goods can only be transported in fixed lot-sizes, then the transportation cost assumes a step-function form,

which adds an extra level of complexity to the facility location problem. In “Approximation Algorithms for Problems Combining Facility Location and Network Design,” R. Ravi and A. Sinha study the optimization of the sum of facility location and network design costs in such situations. They provide heuristics with provable bounds on the solution quality for this problem, relying on simple lower bounds arising from the individual problems of locating facilities and transportation. They also study a version of the problem where there is a bound on the number of facilities that may be opened. While ignoring the effects of the two components of the problem on each other can lead to high costs, this research shows that solutions to the two subproblems can be effectively combined to obtain a close-to-optimal solution for the overall problem.

Multiproduct Pricing Using Online Consumer Preference Data

Developed by General Motors (GM), the Auto Choice Advisor website (<http://www.autochoiceadvisor.com>) recommends vehicles to consumers based on their requirements and budget constraints. Through the website, GM has access to large quantities of data that reflect consumer preferences. In “A Nonparametric Approach to Multiproduct Pricing,” P. Rusmevichientong, B. Van Roy, and P. W. Glynn developed an approach to multiproduct pricing that leverages the availability of such data. They consider models of consumer purchasing behavior, each of which relates observed data on a consumer’s requirements and budget constraint to subsequent purchasing tendencies. They develop a heuristic that optimizes product prices with respect to a sample of consumer data. When applied to the dataset from the Auto Choice Advisor website, their approach provides insights into the current pricing policy at GM and suggests enhancements that may lead to a more effective pricing strategy.

Can You Improve the Performance of Your Optimization Procedure?

Solving difficult optimization problems not only requires many hours of research, development, and implementation to create an effective solution method but also an incredible amount of time to make sure that the procedure performs at the highest possible level. The process by which computer implementations of solution procedures are configured to maximize their performance is tedious and time consuming. In “Fine-Tuning of Algorithms Using Fractional Experimental Designs and Local Search,” B. Adenso-Díaz and M. Laguna describe the development of a tool whose goal is to help operations researchers and practitioners ease the pain of improving the performance of their computer implementations. The tool, referred to as CALIBRA, employs concepts from statistical analysis and search strategies to identify configurations through the adjustment of input parameters. CALIBRA does not intend to substitute for human intuition and expertise during the fine-tuning

process, but rather it attempts to make the search for the “best configuration” a relatively painless experience based on scientific methods. The authors believe that automatic fine-tuning is at an infant stage of development. However, CALIBRA provides an initial step that hopefully motivates future explorations and developments in this area.

A Reliable COMPASS for Optimizing Simulated Systems

Computer simulation is a standard engineering tool for improving and optimizing system design and performance. Stochastic computer simulations incorporate probability models to represent the uncertainty in a system and then generate samples from these models to drive the simulation and estimate the system’s performance. “Optimization” in this setting usually means maximizing or minimizing some measure of long-run average performance. However, due to sampling, the performance that can only be estimated is subject to statistical error. As a result, convergence of optimization algorithms for stochastic simulation problems with discrete decision variables (e.g., number of machines in a work center, staffing levels per time period in a call center, order-up-to levels in a supply chain) typically requires that all feasible solutions are simulated, and simulated infinitely often; this is unrealistic for practical problems with large numbers of feasible solutions. In “Discrete Optimization via Simulation Using COMPASS,” L. J. Hong and B. L. Nelson describe convergent optimization via most-promising area stochastic search (COMPASS), an approach that is simple and easy to understand, is probably convergent, but requires simulating only a small fraction of the available solutions (even when there are infinitely many feasible solutions).

Scheduling Vehicles in an Urban Transit System

One of the main problems that must be solved by urban transit authorities is assigning vehicles to sets of trips (the “duties” of the vehicles). In the special case where all the vehicles start from the same depot, this problem can be modeled as a minimum cost flow problem and solved easily. In practice, however, every city has several depots, and it is difficult to compute an optimal solution of the resulting model, called the MDVSP (multiple depot vehicle scheduling problem). The MDVSP is similar to other important problems in the transportation area, for instance the crew scheduling (or crew pairing) problem in airline transportation. In “A Branch-and-Cut Algorithm for the Multiple Depot Vehicle Scheduling Problem,” A. Hadjar, O. Marcotte, and F. Soumis use several techniques of integer and linear programming (column generation, variable fixing, and cutting planes) to solve large instances of the MDVSP. Their article contains reports of extensive computational experiments using randomly generated data sets, real-life data sets, and data sets found in the literature.

Addressing Uncertainty on the Uncertainty with Robust Optimization in Inventory Management

Real-life inventory problems face uncertainty in customer demand, which is traditionally addressed by assuming specific probability distributions. But these assumptions are only educated guesses of what tomorrow's demand might be like, and mistakes can be costly for the decision maker. Even if the assumptions are correct, stochastic formulations quickly become intractable in dynamic settings. On the other hand, robust optimization has been widely used to address data uncertainty in mathematical programming problems, with applications ranging from portfolio selection to network flows, and yields tractable models. In "A Robust Optimization Approach to Inventory Theory," D. Bertsimas and A. Thiele were motivated by these results to develop a new framework for dynamic optimization under uncertainty. They describe random demands by uncertain parameters belonging in uncertainty sets rather than stochastic variables obeying probabilistic distributions, and as a result obtain linear and mixed integer programming problems as robust, tractable counterparts of the stochastic inventory problem. They derive theoretical insights into the structure of the optimal policy for single stations, series systems, tree networks, and more general supply chains.

Offering Discounts and Allocating Inventory to Multiple Customer Classes

Many wholesalers face the problem of determining which customers to service when inventories run low and demand exceeds supply. One approach they take is to offer discounts for delayed fulfillment to some of their customers. Online and catalog businesses face similar problems as they need to determine the availability to ship on a given day. In "Dynamic Pricing Through Discounts for Optimizing

Multiple-Class Demand Fulfillment," Q. Ding, P. Kouvelis, and J. M. Milner investigate how a firm should determine the discounts and which customers to serve. They advocate the use of combined rationing and discounting policies for environments in which the customer does not see real time inventory levels. Dynamic programming and optimization approaches can be used to determine the optimal inventory levels and customer class based discounts. They show that while discounts are useful in maintaining satisfactory fill rates for the delayed customers, profits are primarily influenced by the inventory and allocation policy followed by the firm. Under the policies, suppliers increase profits and reduce inventory levels, with the inventoried goods primarily reserved for the more profitable and time sensitive customers who request "same day" service. The cost sensitive demand is delayed for "next day" delivery, with the waiting customers compensated for their patience via appropriate discounts. The policies expand the inventory-service efficient frontier, leading to a win-win situation for suppliers and customers in these environments.

Linear Multicommodity Flow

The linear multicommodity flow problem (LMFP) is a standard problem in transportation and telecommunication. A standard way to solve those problems is to resort to Lagrangian relaxation and solve the Lagrangian dual problem via some nondifferentiable optimization technique. In "Solving Large-Scale Linear Multicommodity Flow Problems with an Active Set Strategy and Proximal-ACCPM," F. Babonneau, O. du Merle, and J.-P. Vial exploit the observation that the number of congested arcs is usually limited. The authors use an active set strategy to focus on the congested or nearly congested arcs. The active set strategy is embedded in a recent enhancement of the analytic center cutting plane method. The new approach solves huge instances of LMFP.