



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(3):ii-iv. <https://doi.org/10.1287/opre.1060.0316>

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

How Good Is Your Portfolio Strategy?

Solving for optimal portfolio strategies is a difficult task in many of today's financial models. Complex security price dynamics, portfolio constraints, parameter uncertainty, and trading frictions all combine to create a very difficult stochastic control problem. These problems cannot be solved to optimality, and so instead we have to make do with suboptimal strategies. Unfortunately, it is usually difficult to tell how good these suboptimal strategies are. In particular, how much better off would you be if you could use the unknown optimal strategy? In "Evaluating Portfolio Policies: A Duality Approach," M. B. Haugh, L. Kogan, and J. Wang show how these suboptimal strategies may be used to create lower and upper bounds on the performance of the unknown optimal policy. As a result, they can also determine how far a given suboptimal strategy is from optimality.

Staffing and Routing in Telephone Call Centers

Telephone call centers are a rich source of problems for operations researchers, with two generic problems dominating the OR literature: the *call center staffing problem*, where personnel levels are determined for each of several multiskilled agent pools; and the *dynamic call routing problem*, where incoming calls are assigned to agents on a real-time basis. A. Bassamboo, J. M. Harrison, and A. Zeevi describe an integrated approach to the staffing and routing problems in "Design and Control of a Large Call Center: Asymptotic Analysis of an LP-Based Method." In their approach, optimal staffing reduces to a stochastic programming problem—one that is computationally tractable for systems of realistic size—and dynamic call routing is accomplished by frequently re-solving a linear program of moderate size, using real-time estimates of average demand rates to vary server allocations as time passes. The method is shown to be nearly optimal in a parameter regime of practical interest.

New Optimality Conditions in Nonconvex Optimization and Their Uses

Algorithms in nonlinear optimization are founded on conditions describing the existence of Lagrange multipliers at some point of stationarity or optimality. In integer and combinatorial optimization this type of condition was previously largely unavailable. For such problems, therefore,

there has historically been a gap between the application of multiplier-based solution strategies, such as Lagrangian heuristics, column generation and core problems, and optimality conditions founded on Lagrangian multiplier theory. In "Global Optimality Conditions for Discrete and Nonconvex Optimization—With Applications to Lagrangian Heuristics and Column Generation" by T. Larsson and M. Patriksson, this gap is closed. Among its chief consequences, the new optimality conditions may serve both as the basis of the analysis of already existing computational approaches in integer programming and as sources of inspiration for new ones. In particular, the authors provide several proposals for the construction of column generation algorithms for integer programs.

Stabilizing and Regularizing Column Generation Behavior

Column generation users experience instability features as the size and the difficulty level of the problems increase. For example, primal degeneracy is well known to cause slow convergence and tail effect. "Dual-Optimal Inequalities for Stabilized Column Generation" by H. Ben Amor, J. Desrosiers, and J. M. Valério de Carvalho provides a general framework on the use of relevant dual information as an alternative tool for stabilizing or regularizing column generation behavior.

A Simple and Efficient Scheduling Rule for Large-Scale Dynamic Systems

Most traditional research on scheduling problems has focused on static or deterministic problems in which all jobs are available for processing at the same time and the job parameters are known before processing starts. In practice, however, jobs typically arrive dynamically to the system and job parameters are not known in advance, e.g., the exact job processing time is not known until processing is completed. In addition, most solution techniques developed have focused on small size scheduling problems, whereas in practice it is common to have scheduling problems with thousands of jobs that need to be processed on one or more machines. In "On the Asymptotic Optimality of a Simple On-line Algorithm for the Stochastic Single-Machine Weighted Completion Time Problem and Its Extensions," M. C. Chou, H. Liu, M. Queyranne, and D. Simchi-Levi propose a simple and efficient algorithm to minimize the average job completion time for large-scale dynamic systems. Such a system may represent a flow shop or an open

shop scheduling problem where jobs arrive dynamically to the system, the information about future arrival of a job and its parameters is not available until the job arrives to the system, and the exact job processing time is not known until the processing completes.

Comparison of Long-Run Performance

Simulation experiments are often used to estimate the performance of systems, e.g., production, financial, service, and communication, whose behavior is at least partially determined by uncontrollable events such as raw material supplies, interest rates, customer orders, and user traffic, respectively. To compare alternative system designs, and in particular to find the best design, with a statistical guarantee of being right is difficult when long-run (“steady-state”) performance is what matters. In “On the Asymptotic Validity of Fully Sequential Selection Procedures for Steady-State Simulation,” S.-H. Kim and B. L. Nelson provide the first theoretical framework within which valid statistical procedures can be constructed for practical simulation problems, and they design two such procedures.

Rare Event Simulation of Markov Chains

Rare event analysis is critical in many practical settings including finance, insurance, telecommunications, and reliability systems. Performance evaluation of rare events via naïve simulation is difficult because it may take enormous computational effort to generate rare events. In the last 20 years, significant research has focused on developing importance sampling simulation techniques to quickly estimate extremely small probabilities. When successful, this could reduce computational effort by a factor of thousands. However, the successes of the proposed techniques have been limited to special classes of Markov chains such as those used to model highly reliable systems or those that involve random walks without boundaries. In recent years, exciting adaptive importance sampling techniques have been proposed to quickly simulate general Markov chains. In “Adaptive Importance Sampling Technique for Markov Chains Using Stochastic Approximation,” T. P. I. Ahamed, V. S. Borkar, and S. Juneja propose a new stochastic approximation based adaptive importance sampling technique to efficiently estimate many commonly encountered performance measures associated with Markov chains. Their approach differs from the existing adaptive approaches that are based on generating complete sample paths to rare events. Through extensive experiments, they show that the proposed algorithm significantly outperforms the existing ones in estimating rare events.

Practical Bounds and Schedules for Multiproduct Single-Server Systems

Manufacturing environments that produce a number of different products in a make-to-order fashion on a single

machine, where setup activities are necessary when making switches of product type, are relatively common in industry. In general, the arrival and processing rates are likely to differ across product types. Because of these differences, production scheduling will have a significant impact on production costs. In “Multiproduct Systems with Both Setup Times and Costs: Fluid Bounds and Schedules,” W.-M. Lan and T. L. Olsen give performance bounds for evaluating scheduling heuristics for such systems and show how these bounds lead to good heuristic schedules. These bounds and heuristics are based on a fluid model analysis that yields simple yet robust results.

Lot Scheduling Problems with Safety Stock

The economic lot scheduling problem is an important building block in the operation management literature. In “The Multiple-Family ELSP with Safety Stocks,” S. M. Karalli and A. D. Flowers address an important extension to this problem; namely the explicit treatment of safety stock costs in the solution procedures. They use a basic period approach combined with family and item multipliers to attempt to minimize setup plus inventory carrying costs, with the latter including safety stocks.

Systems for Reverse Logistics

Managers and scholars have begun the important but challenging task of extending logistics systems to encompass return flows. This concern with reverse logistics is driven partly by environmental issues but also by economics—a return option provides real value to customers, and the recovered products provide real value to the firm. In “Optimal Policy for a Multiechelon Inventory System with Remanufacturing,” G. A. DeCroix explores the implications for inventory management of returns in a serial system. In general, the optimal policy for managing the system is a natural combination of the optimal policies for two related systems—a serial system without returns, and a single-facility system with returns. However, depending on the form in which returned items are recovered, e.g., raw components vs. partially manufactured end products, this policy structure may require some modifications in how inventory is counted and in what actions are allowed.

A Supply-Demand Network Equilibrium Model for Cost Reduction

A supply-demand network is a collection of sources of supply where products are made and then distributed through some intermediaries, and finally to a collection of points of demand/consumption. An important problem in supply chain management is the coordination of product and material flow among locations scattered in the supply chain. Such a problem typically involves transporting products located at various facilities to geographically dispersed

retailers at minimum cost or with shortest delay. In “A Multiproduct, Multicriterion Supply-Demand Network Equilibrium Model,” T. C. E. Cheng and Y. N. Wu present an analytical framework based on the concept of network equilibrium to attain optimal performance for a supply-demand network involving multiple products and multiple criteria cost functions.

Integrating Order Assignment, Processing, and Delivery Decisions in a Supply Chain

Time-sensitive product manufacturers often operate in a make-to-order mode such that finished products are delivered to customers soon after they are produced and little finished product inventory is kept. When making production and distribution decisions, such manufacturers have to consider trade-offs between responsiveness to their customers and total production and distribution cost. Many existing models have considered such trade-offs at a strategic or tactical level. Those models often involve inventory decisions in addition to production and distribution. However, detailed scheduling level decisions are sometimes more important and relevant in optimizing such trade-offs for a time-sensitive product manufacturer. In “Order Assignment and Scheduling in a Supply Chain,” Z.-L. Chen and G. Pundoor consider a scheduling model that optimizes such trade-offs. The model integrates three types of decisions: assigning orders to plants, scheduling of order processing operations at each plant, and scheduling delivery to a central distribution center.

Solving the Dial-a-Ride Problem by Branch-and-Cut

In recent years, transportation-on-demand systems have become increasingly popular in many countries. With the aging of the population and the trend toward the development of ambulatory health care services, more and more people rely on dial-a-ride transportation systems provided by local authorities. In the dial-a-ride problem, users formulate requests for transportation from a specific origin to a specific destination. Transportation is carried out by vehicles providing a shared service. The problem consists of designing a set of minimum-cost vehicle routes satisfying capacity, duration, time window, pairing, precedence, and ride-time constraints.

In “A Branch-and-Cut Algorithm for the Dial-a-Ride Problem,” J.-F. Cordeau applies branch-and-cut to solve the

problem efficiently. In this vehicle routing problem, precedence constraints highly influence the structure of the problem and give rise to several families of strong inequalities as well as many problem reduction techniques.

The Irregular Packing Problem: Increasing Profits by Reducing Waste

The process of cutting shapes from a sheet, or multiple sheets, of material is a common requirement within many important manufacturing industries. In such industries, costs are often directly related to the labor and material involved within the cutting operations. Generally, a reduction in labor would have an adverse effect on productivity and, therefore, cost reduction is restricted to the more efficient utilization of raw material. This is most aptly demonstrated through the example of a mass production industry where a template layout of shapes is cut repeatedly. If just small improvements can be made in material utilization, then this can translate into very significant and accumulative cost savings. In “A New Bottom-Left-Fill Heuristic Algorithm for the Two-Dimensional Irregular Packing Problem,” E. Burke, R. Hellier, G. Kendall, and G. Whitwell propose a new algorithm for achieving high quality arrangements of irregular shapes on a given sheet. Their proposed approach produces 25 best-known solutions on 26 well-established benchmark data and does so in faster time when compared to previously published algorithms. Furthermore, the authors’ algorithm is capable of packing irregular shapes that may contain any number of lines, arcs, and holes and this, in itself, represents a significant improvement over the previous state of the art. Finally, the authors introduce and set initial results for 10 new benchmark instances, which contain a variety of shapes with arcs and holes, with the hope that other researchers will be encouraged to develop algorithms that are capable of packing such shapes.

Verifying the Increasing Generalized Failure Rate (IGFR) Property

Frequently, research on pricing or supply chain contracting problems assumes a probability distribution with an increasing generalized failure rate (IGFR). While it is known that this class of distributions is large, proving that a specific distribution falls in this class can be cumbersome. In “A Note on Probability Distributions with Increasing Generalized Failure Rates,” M. A. Lariviere provides several alternative characterizations that simplify the process of verifying the IGFR property.