Operations Research

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

In This Issue

To cite this article:
(2002) In This Issue. Operations Research 50(2):ii-iv. https://doi.org/10.1287/opre.50.2.0.431

Trimming Paper Companies' Costs with Decision Support

The production of paper is a highly capital intensive manufacturing process. In order to have a strong return on investment and survive in a highly competitive environment, paper companies must be responsive to customer demand and ensure short production and distribution lead times by using their resources efficiently. A key factor in meeting these goals is to develop an effective strategy of production and distribution scheduling. In "Scheduling Solutions for the Paper Industry," P. Keskinocak, F. Wu, R. Goodwin, S. Murthy, R. Akkiraju, S. Kumaran, and A. Derebail describe a decision support system for paper mill scheduling, which considers the discrepancies between multiple stages of manufacturing and distribution, as well as multiple, and often conflicting, objectives. The agent-based architecture of the system allows alternative solution methods to cooperate with one another and with the schedulers to generate multiple high quality solutions.

Algorithms for Solving Routing Problems in Telecommunication Networks

The routing of traffic requirements is one of the main problems that telecommunications companies have to solve routinely. Each requirement is specified by a pair of nodes between which a given amount of traffic has to be routed through the network. Given the network link capacities, the objective is to find a routing of the traffic requirements that minimizes the cost and is such that the accumulated flow on each link does not exceed its capacity. This problem can be modeled as a nonoriented multicommodity flow problem for which the authors study, propose, and compare specialized solution methods based on simplex and interior-point techniques. In "Simplex and Interior Point Specialized Algorithms for Solving Nonoriented Multicommodity Flow Problems," P. Chardaire and A. Lisser report results on test instances derived from the France Telecom Paris district transmission network. In particular, their numerical experiments demonstrate that solutions to such typical instances are obtained very efficiently by algorithms based on price decomposition techniques.

Using Survival Analysis to Build Credit Scoring Models

Credit scoring is one of the most widely used applications of operations research and statistics in business today. There is currently a practical need for tools that will allow credit scoring to develop from its present role in predicting default, to a new focus on its ability to predict profit. The profitability of an account is strongly related to the length of time the consumer stays with the lender; therefore, the requirement is to change to a time-based measurement, and thus a change in the techniques being used. So far there have been only preliminary investigations that show that survival analysis models, and in particular proportional hazards models, are competitive with the current industry standard technique, logistic regression. In "Survival Analysis Methods for Personal Loan Data," M. Stepanova and L. Thomas develop three more facets of the proportional hazards approach that will be useful in practice. These include using survival analysis rather than logistic regression to segment continuous variables in order to deal with nonlinearities, investigating appropriate residual diagnostics of the model's fit, and adopting the idea of time-dependent covariates to overcome the assumption of proportional hazards that the same group are most at risk at all times.

A New Criterion for Burn-In

The determination of the time at which a device is burned in is crucial in many applications. In "A Criterion for Burn-In That Balances Mean Residual Life and Residual Variance," H. W. Block, T. H. Savits, and H. Singh attempt to determine such a time using the concepts of mean residual life and residual variance. Generally for lifetimes for which burn-in is appropriate, e.g., distributions with bathtub-shaped failure rates, the mean residual life increases, then decreases. The residual variance also increases, then decreases. A general result is developed for studying these types of behaviors. Finally, the residual coefficient of variation, which combines the mean residual life and the residual variance, is analyzed, and it is shown how an optimal burn-in time can be determined using this concept. The optimal time is compared to other optima obtained using other criteria.

Asset Allocation in Deregulated Electricity Markets

Though now in its third year of deregulation, California's electricity market continues to exhibit extremely volatile prices and even occasional blackouts. Moreover, such problems are not unique to California. Almost without exception, participants in deregulating electricity markets around the world share a similar concern: How should they manage these new risks? For owners of generation assets, the problem is compounded by the numerous operating constraints of the assets themselves and the similar, albeit smaller, underlying uncertainty in fuel prices. By extending recent advances in real options pricing techniques and using financial models tailored to observed electricity and fuel price behavior, the authors have developed a new valuation model to address such concerns. In "Short-Term Generation Asset Valuation: A Real Options Approach," C. L. Tseng and G. Barz investigate the relationship between the value of a power plant and its operational flexibility, consequently helping asset owners who are looking to capitalize the option value of their physically-constrained assets. Because

it incorporates asset constraints, this research also provides insight into the optimal operational policy for a power plant under such price uncertainty.

Optimal Search for an Agent of Unknown Aims

When searching for a person who is walking in a known area, it is sometimes known whether he wants to be found (e.g., a lost hiker) or whether he does not (e.g., an escaped prisoner). In the former case, the theory of rendezvous search might be useful; in the latter, the theory of search games. However until recently there was no satisfactory theory to deal with the scenario in which the aims of the "target" agent are only known probabilistically, as in the case of a lost teenager who might in fact be a "runaway." In "Searching for an Agent Who May or May Not Want to Be Found," S. Alpern and S. Gal develop just such a theory. For a given search area, they consider a family of rendezvous-evasion games parameterized by a probability c that the agent is a "cooperator" who wants to be found. This family connects the rendezvous problem (c=1) by a continuous path to the search game (c=0), thus establishing a connection between two previously distinct areas of search theory. The authors establish complete solutions (for all values of c) when the search area is a finite tree and when the two players are placed a known distance apart on a line (with a fast searcher and a slow agent).

Improving Order Allocation in the Paper Industry

A common problem in any firm with many somewhat similar machines or production facilities is that of deciding which machine to use for each order or job as it arrives. A complication encountered in the paper industry is that the efficiency with which a particular machine can process a particular order depends upon what other orders are assigned to that machine. In the paper industry, this assignment of orders to machines had always been done manually because standard optimization procedures were too time consuming to solve the complete problem. In "Order Allocation for Stock Cutting in the Paper Industry," S. Menon and L. Schrage show that, by developing some computationally cheap bounds on the cost of an allocation and then exploiting the decomposition structure of the full problem, one can quickly get solutions to the full problem encountered in the paper industry that are typically within 1 percent or 2 percent of optimum.

Designing Telecom Networks with Several Capacity Options

This paper addresses the design of telecom networks when several options are available for installing capacity on the links. Under each option a definite capacity can be installed on the link at a given price. Most previous work on network design has assumed that capacity can be installed in multiples of one base unit (e.g., 45 megabits/second for a T3-based network). This is usually the case when telecom carriers design their own transport networks using high capacity transmission facilities of a fixed size. However, the carriers offer several capacity options to their customers. For example, the fractional T1 services typically offer bandwidth to customers in units of 1, 2, 4, 6, 8, 12, and 24 DSOs (one DSO is equivalent to a 64 kilobits/second channel). The pricing of these services usually has a built-in volume discount, i.e., die per unit cost of capacity decreases, as one moves from lower to higher capacity services. This adds a new dimension of complexity to cost/capacity trade-off in a network design algorithm. In "Design of Capacitated Multicommodity Networks with Multiple Facilities," Y. Agarwal proposes a heuristic approach that identifies a relatively easy subproblem and solves it repeatedly to improve the network design. Computational results on a variety of test problems suggest that this approach produces solutions within 5 percent of optimum in most cases.

Analysis of a Least Recently Used Cache Management Policy for Web Browsers

With the explosive growth of Internet enabled business, companies are rushing to build more attractive websites, often with rich multimedia content. Internet users need to download more information while browsing the web. There have been many cases of users becoming frustrated by long delays in loading web pages. One effective way to reduce the wait is through *caching* – the practice of storing certain documents on the user's hard disk. Software vendors have incorporated various caching policies in the web browser. In "Analysis of a Least Recently Used Cache Management Policy for Web Browsers," V. Mookerjee and Y. Tan analyze an LRU (Least Recently Used) policy for cache management in a web browser. They derive analytical expressions to evaluate performance and present an accurate and easy to compute approximation. Browser designers can apply the results of the study to adopt a suitable cache management policy. New versions of browsers can also include a simple cost-benefit analysis to help users to optimize the cache size. The study may also motivate other researchers to analyze more complex caching policies using the techniques presented here.

Bounds on Optimal Pricing

A central question in financial economics is to find the price of a derivative security given information on the underlying asset. Under the assumption that the price of the underlying asset follows a geometric Brownian motion and using the no-arbitrage assumption, the Black–Scholes formula provides an explicit and insightful answer to this question. This research makes an explicit connection between asset pricing in finance and linear and semidefinite optimization by providing best possible bounds on prices of options just from the assumption of no-arbitrage, and without assuming a specific model for the underlying price dynamics. In "On the Relation Between Option and Stock Prices: A Convex Optimization Approach," D. Bertsimas and I. Popescu find that semidefinite optimization methods, duality, and complexity theory shed new light on the relation

of option and stock prices. For the single stock problem, given moments of the prices of the underlying assets, they show that one can find best possible bounds on option prices with general payoff functions efficiently. Conversely, given observable option prices, they provide best possible bounds on moments of the prices of the underlying assets by solving linear optimization problems. For options that are affected by multiple stocks, they find nonoptimal bounds using convex optimization methods. However, they show that it is NP-hard to find best possible bounds in multiple dimensions. The results are extended to incorporate transactions costs.

Optimal Service Capacity with Upgrades

Ever wondered why, when you show up at the airport's car rental desk to pick up that compact car, they so often seem to upgrade you to a mid-sized one? Clearly the car rental company is making a good decision to upgrade you rather than to turn you out into the street without a car. But this upgrade policy complicates the capacity decision; the firm must decide how many cars of each type to keep on-hand, given the option to upgrade customers when they arrive. In addition, correlation in the demand for car sizes may affect this capacity decision. In "Flexible Service Capacity: Optimal Investment and the Impact of Demand Correlation," S. Netessine, G. Dobson, and R. Shumsky examine capacity decisions when customer upgrades are possible. The results also apply to airline seats and finding optimal inventory levels of flexible goods such as the CPUs for personal computers. The authors look at a single-period model in which upgrades are only allowed to the next-highest class. They show that if demand follows a multivariate normal distribution, the optimal capacity would shift from mid-sized to compact cars as the correlation of demand between the two increases. In addition, the capacity for other sizes, both larger and smaller, would alternately rise and fall. This result resolves conjectures, made by numerous authors, about how correlation affects optimal capacity.

Estimating Time-Dependent Attrition Coefficients for Lanchester Square Law

The techniques of inverse problems were applied in many different areas of research. Recently, a more general approach (i.e., a nonparametric conjugate gradient method (NP-CGM)) to such problems has been developed in the engineering literature. It is specifically tailored to handle the inverse problem of using system response data to estimate time-varying coefficients in the differential equations that describe system dynamics. In "An Inverse Problem of the Lanchester Square Law in Estimating Time-Dependent Attrition Coefficients," H. Chen uses it to estimate time-varying attrition coefficients in Lanchester's square law by using observed data on some or all of the battle's strength histories. She uses hypothetical strength histories and reinforcement schedules that are known without error at several points in time to illustrate the method. In contrast to the method used in Hartley III and Helmbold, the NP-CGM does not require any assumption about the constancy or piecewise constancy of the attrition coefficients. Nor does it employ a particular functional form for the time-varying attrition coefficient, characterized by a finite number of adjustable parameters, as would be needed to apply the parametric conjugate gradient method (P-CGM).