Operations Research

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

In This Issue

To cite this article:
(1996) In This Issue. Operations Research 44(5):662-664. https://doi.org/10.1287/opre.44.5.662

IN THIS ISSUE

Settling a $94 Billion Debt Entanglement

In August 1982, the default of a major trader led to the collapse of Kuwait's al-Manakh stock market, an unofficial and highly speculative equity market. Post-crash debts totaled US $94 billion, or more than four times Kuwait's Gross Domestic Product. The debacle threatened the stability of the economy, securities lost much of their peak values, and the real estate market plummeted. Traders who became insolvent could pay only a fraction of their debts. Courts could not settle traders' debts, however, due to their simultaneous entanglements with each other. To avoid a legal nightmare, the Finance Minister called on the authors to develop quantitative method(s) to solve the post-crash entanglements. Abdelghani Elimam, Maurice Girgis, and Samir Kotob designed a linear programming model that used traders' receivables, payables, and assets to identify solvent from insolvent traders, untangle traders' debts from each other, and decide how each insolvent trader will pay his creditors by type of asset. The courts adopted the results of the models. In the final press briefing on the crisis, the Finance Minister praised the models as effective, equitable, and robust.

Forest Fire Managers Use OR to Evaluate Airtanker Home Basing Strategies

The Ontario Ministry of Natural Resources uses amphibious airtankers to help control forest fires. The aircraft are based at a small number of airports that serve as "home bases" from which the aircraft are deployed to a larger set of airports that serve as initial attack bases for fire fighting operations. Fire managers must develop hierarchical strategies that stipulate where the "home base" for each airtanker will be home based and how they are to be deployed each day given the highly uncertain demand for airtankers that varies over time and space. In "Basing Airtankers for Forest Fire Control in Ontario," James MacLellan and David Martell describe how they helped Ontario's fire managers use daily forecast weather and fire occurrence predictions to derive subjective airtanker deployment rules. These rules specified how many airtankers were to be deployed at each initial attack base each day of the fire season, and incorporated those rules in a mathematical programming model. This model was used to identify a home basing strategy that would minimize the average annual cost of satisfying daily airtanker deployment demands. Their analysis provided valuable insight that was used to develop a new airtanker home basing strategy that has been implemented in Ontario.

Cocaine: How Can It Be Controlled?

The production, distribution, and consumption of cocaine impose significant costs on American society. Government at all levels has responded by spending approximately $13 billion per year on a variety of cocaine control programs including source country control, interdiction, domestic enforcement, and treatment, among others. Despite such substantial spending, there has been essentially no systematic effort to quantify the relative cost-effectiveness of these diverse programs. In "Enforcement of Treatment? Modeling the Relative Efficacy of Alternatives of Controlling Cocaine," C. P. Rydell, J. Caulkins, and S. Everingham construct a model that provides such estimates. They find that treatment is the most cost-effective, followed in order by domestic enforcement, interdiction, and source country control. This result should reinforce often-heard intuitive and emotional pleas for a shift in drug control efforts away from the current emphasis on supply-side programs toward providing additional treatment.

Drug Therapies: Establishing Control Policies

An important problem in drug therapy is to learn how patients respond to medication while concurrently treating the medical condition. This problem can be formulated into the partially observed Markov decision process framework, and the application to this area has up to this point been dominated by naive feedback control policies, which neither hedge against estimation errors nor account for the value of active information seeking. In "Comparison of Some Suboptimal Control Policies in Medical Drug Therapy," C. Hu, W. S. Lovejoy, and S. L. Shafer define some computationally tractable policies for adaptive control of drug levels, and identify an efficient control scheme for drug delivery via compared simulated patient applications. In their simulation study, a suboptimal control policy that hedges against estimation errors is shown to be superior to naive feedback control. The authors conclude that hedging against estimation errors is an important component of optimal control of drug therapy, but that active information probing is not cost effective in this setting.

Routing Hazardous Materials

In "On the Analysis of Two New Models for Transporting Hazardous Materials," Honghua Jin, Rajan Batta, and Mark H. Karwan analyze two models for hazardous materials transportation that are based on a fundamentally different concept emanating from viewing hazmat routing as a probabilistic experiment. The first model focuses on finding a path that minimizes the expected total consequence when the scenario ends after a threshold number of accidents and a finite number of shipments are involved. The second model has the objective of minimizing the expected consequence per trip of an infinite number of shipments under a similar threshold on the number of permissible

accidents. Constraints on total accident probability and total expected consequences are applied to both models, and exact solution methods are proposed and computationally tested.

Insights at High Service Levels

Effective allocation of production capacity among multiple items is a difficult problem. In "Allocating Production Capacity among Multiple Products," Paul Glasserman provides a solution that becomes optimal at high service levels. The analysis proceeds from the modeling assumption that a fixed proportion of overall capacity can be dedicated exclusively to the production of each item. For each allocation of capacity, production of each item follows a base-stock policy. The objective is to minimize holding and backorder costs, or to minimize holding costs subject to a service-level constraint. Numerical results indicate that the asymptotically optimal allocation rules perform very well, even far from the asymptotic regime.

Finding the Groups Quickly

The cell formation problem in group technology is a hard combinatorial optimization problem. The goal in "On the Characterization and Measure of Machine Cells in Group Technology," is to find an efficient algorithm that can find a good approximation solution quickly with respect to a given measure. The measures used in this class of problems are also critical in determining the quality of a solution. The approach presented in this paper is also able to identify the independent machine cells in a large system easily. Efficient criterion to break up a larger machine cell into suitable smaller machine cells is included. This approach has been used to tackle problems in seemingly unrelated industries over the last few years: to quickly identify the document groups and customer groups to minimize telecommunication cost in an automatic information retrieval system; and to cluster the job orders in the scheduling of a production system of a liquid crystal display manufacturer.

A Simple but Powerful Approach for Analyzing Performability

Connectedness reliability is one of the fundamental problems in the area of performability analysis as it applies to communication and transportation networks, and electronic and VLSI circuitry. Performability analysis gives network designers a tool for assessing the strength and robustness of a network. Performability problems are among the most intractable combinatorial problems considered in practice, even when good approximations—rather than exact values—are desired. In "The Delta-Wye Approximation Procedure for Two-terminal Reliability," Manoj K. Chari, Thomas A. Feo, and J. Scott Provan introduce a unique and elegant new way to approximate performance measures by combining a simple but effective network reduction procedure with a tight "local" probability approximation analysis. Empirical studies indicate that this procedure is not just "better" than previous methods, but better by an order of magnitude. The network reduction procedure, furthermore, is generic and can be applied to many other kinds of problems, including more sophisticated performance measures such as throughput, transmission delay, and reachability.

How Are Big Delays Related to Big Workloads?

In queueing applications, we may be interested in the waiting time (time from arrival until starting service), sojourn time (waiting time plus service time), or the workload (remaining work of all customers in the system at an arbitrary time). We often are interested in probability distributions as well as means. In their paper, "Exponential Approximations for Tail Probabilities in Queues, II: Sojourn Time and Workload," Joseph Abate, Gagan Choudhury, and Ward Whitt show that exponential asymptotics for the tail probabilities of these three distributions are closely related, so that it is easy to go from one exponential approximation to another.

SPT Strikes Again

Recently, a new special class of job shops called *reentrant shops* came into prominence. The basic feature of a reentrant shop is that jobs visit a certain machine, called the *hub*, many times. The last decade has seen a growth of reentrant shops, especially in the high-tech industry. For example, in VLSI (Very Large Scale Integrated) circuit wafer fabrication facilities, wafers—thin silicon slices on which semiconductor devices are built—have to be processed more than ten times on a photolithography workstation. Between two successive visits to the photolithography workstation, the wafers are processed on the other machines. In "Mean Flow Time Minimization in Reentrant Job Shops with a Hub," W. Kubiak, S. X. C. Lou, and Y. M. Wang study reentrant shops with the objective to minimize mean flow time. Because of the model's computational intractability, they study it under assumptions that are satisfied in certain real world situations. The assumptions allow the authors to reduce the whole shop to a single machine reentrant job shop with setups between different entries of the same job. The authors then show that there are optimal schedules with the shortest processing time (SPT) job order and they propose a dynamic programming algorithm to find such schedules. The authors analyze two special cases with the restriction that there are only two entries to the hub and offer more insights into the structure of optimal solutions.

Two-stage Production Schedules with Waiting Times

In production scheduling on two different machines, the Flow Shop and Job Shop models are the major paradigms in theory and in practice. However, these production schemes do not apply to several practical cases where each job has only two operations, but it is necessary to leave a minimum time interval between the end of the processing

of the first operation and the beginning of the second one. Consider, for example, a steel industry where big hot slabs are produced from the fused metal and later, when the temperature has dropped, the slabs have to be cut to appropriate sizes. The cooling phase determines a minimum waiting time between the two operations. The classical two-machines (Flow Shop and Job Shop) problems are considered by Mauro Dell'Amico in "Shop Problems with Two Machines and Time Lags." The computational complexity of these problems is studied, lower bound and approximation procedures are given, and the worst-case performance is analyzed. Finally a local search algorithm is developed and the effectiveness of the various procedures is evaluated through extensive computational experiment.

Risk-sensitive Dynamic Medical Treatment Choice via Generalized Utility Function Rollback

Models of medical treatment choice typically involve both uncertainties that resolve at particular points in time—such as surgical success versus failure, or stroke survival versus mortality—and uncertainties which stretch out over time—such as the disabling effects of stroke, or the long-term side effects of drugs. For this reason, Gordon B. Hazen and James M. Pellissier have found it natural to depict such models as diagrams that are hybrids of decision trees and continuous-time Markov chains. The authors call such diagrams *stochastic trees*. The computation of mean quality-adjusted lifetime in a stochastic tree may be performed using a convenient rollback property. However, mean quality-adjusted lifetime can represent only *risk-neutral* preferences. It is desirable to be able to use more general *risk-sensitive* utility functions, not all of which have convenient rollback properties. Hazen and Pellissier show that the practical feasibility of stochastic tree rollback for a general utility function depends on the existence of a convenient *preference summary*. They use these results to illustrate the impact of risk attitude on a patient's decision to undergo a surgical clearing of the carotid arteries.

Departure Processes of Queues

Characterization of departure processes for queueing systems is an important issue in queueing theory. It plays an essential role in queueing network approximation and is often used in conjunction with a decomposition method. In "The Departure Process of the $GI/G/1$ Queue and Its MacLaurin Series," Jian-Qiang Hu uses a novel method to analyze the departure process of the $GI/G/1$ queue. By using this method, he is able to develop a simple recursive procedure for calculating the MacLaurin series of the moments and covariances of the departure process with respect to a parameter in the service time. The results can be used to develop new queueing network approximations; the issue of convergence is also studied.

More Reliable Random Numbers

Many algorithms in operations research and other computational fields involve the generation of (often billions of) random numbers. Linear congruential generators are still widely used for that purpose. However, the relentless increase in available computing power has made those simple generators unreliable for several applications. Other classes of generators have been proposed and studied, including some based on combined linear congruential recurrences. The nonlinear generators are typically much slower than the linear ones. In "Combined Multiple Recursive Random Number Generators," Pierre L'Ecuyer studies the combination of linear congruential recurrences of an order larger than one. He shows that such a combination offers a good way to obtain a reasonably fast and reliable generator, which improves significantly upon the combination of simple linear congruential recurrences. The paper also provides a specific generator with its computer implementation.

Optimizing Routing Policies

In "Adaptive Routing on the Plane," L. Tassiulas presents a new policy for routing a server to service requests arising randomly on a region of the plane. The problem arises naturally in several practical situations including, among others, the dispatching of a repairman of a utility company, failures occurring in a geographical area, or the routing of a supplier that delivers merchandise within a geographical area. Routing policies that maximize the rate of served requests are obtained. The policies are adaptive, that is, they are independent of the statistical parameters of the system. Unlike the previous work on the subject, the results here are obtained for every sample path of the locations of the service requests. Therefore no assumption on the distribution of the demands on the plane is needed and the location of different demands can be arbitrarily depended.