Operations Research

In This Issue

To cite this article:
(1990) In This Issue. Operations Research 38(5):744-931. https://doi.org/10.1287/opre.38.5.744

IN THIS ISSUE

Does OR Address Strategy?

A recent theme of several commentators is that operations research successfully addresses lower level operational problems but not the more strategy-oriented issues that are a central concern of top management in business and government. These commentaries give a misleadingly narrow picture of operations research, which is potentially dangerous for the development of the field. In this OR Forum article, Craig Kirkwood argues that: it is important for us as a profession to address strategic issues, some OR professionals are successfully addressing such issues, and we should give more visibility to this work. If the OR community does not recognize and support the ongoing work on strategic issues, then the practitioners will migrate to other communities. This will mean lost opportunities for OR professionals to work on interesting issues, and less visibility for operations research with corporate and government top level management.

Solving Cartography Problems

Placing labels on a map to maximize legibility and avoid overplot conflicts between labels and between labels and symbols is a tedious process which consumes much of the time required to compile a map by hand. Although many map compilation tasks are automated, label placement is still done by hand or interactively. In fact, label placement for an important class of map features is a combinatorial optimization problem that can be formulated as a multiple choice with thousands of constraints. In "The Solution of Large 0–1 Integer Programming Problems Encountered in Automated Cartography," Steven Zoraster describes a computer program developed to place labels on a type of map widely used in the petroleum industry. The program uses a Lagrangian heuristic to produce solutions to the map label placement problem, and is a new example of the use of mathematical optimization techniques in map design.

Protecting Confidentiality

The U.S. Bureau of the Census must guarantee the anonymity of respondents when publishing tabular data. If the identities of individuals or families could be determined from published data, then the quality of the census would deteriorate because respondents would be less likely to participate or to answer questions honestly. In "Controlled Rounding of Tabular Data," James Kelly, Bruce Golden, Arjang Assad and Edward Baker describe a perturbation technique for providing confidentiality protection. They develop algorithms that combine backtrack search and linear programming techniques for solving the controlled rounding problem associated with three-dimensional tables. Computational results are presented which demonstrate the efficiency of the solution procedures.

Planning Electric Utility Expansion

Until the early 1970s, electric utility capacity planning considered as options supply-side resources whose operation was controlled by the utility, including hydro, oil-fired, gas-fired, coal-fired, and nuclear generation units. However, following the oil embargo of 1973, and with the depletion of fossil fuel and mineral resources, utilities have been considering nondispatchable (not controlled by the utility) resources, such as solar or wind power. Also, integrated resource planning techniques that seek a least-cost mix of supply and demand-side resources, such as conservation and load management, are being applied. In "A Two-Phase Decomposition Approach for Electric Utility Capacity Expansion Planning Including Nondispatchable Technologies," H. D. Sherali and K. Staschus present a capacity planning approach that finds a least-cost mix of dispatchable technologies and nondispatchable resources as well as certain demand-side resources.

Mid-Air Aircraft Refueling

In "Single Aircraft Mid-Air Refueling Using Spherical Distances," Abdulrahman Yamani, Thom Hodgson and Louis Martin-Vega address the problem of a military aircraft that must carry cargo from an origin to some destination base. Mid-air refueling is required by a tanker aircraft that originates from and returns to a third base. The objective is to determine the initial fuel for each aircraft and the location for mid-air refueling to minimize the total fuel consumed by the aircraft. The problem can be an important part of the planning process of military deployments. An efficient solution procedure is developed for the convex programming problem. This is a new application for location problems on a sphere. Given the cost of fuel, and the sheer size of many military deployments, the potential cost savings are significant.

Setting Simulation Run-Lengths

There are two approaches to setting the run-length of a simulation experiment. The first approach preassigns the number of observations to be produced. Because the amount of computer time to generate an observation will typically be random, this implies that the amount of computer time used will be random and, hence, unpredictable. The second approach preassigns the amount of computer time to be expended, in which case the number of observations will be random; this is the more desirable approach when the computer time budget must be kept within a strict bound. The majority of the simulation literature deals with problems that arise in the first setting. This paper, "Bias Properties of Budget Constrained Simulations," by Peter Glynn and Philip Heidelberger is one of the first to analyze the new issues that arise in simulation experiments of the second type, in which the total number of observations collected is random.

Choosing Among Risks

Patients have to choose among different medical treatments with various risks and costs; regulators have to balance risk and costs in prescribing new policies. What criterion should be used in comparing two risk-reduction policies? Some criteria—the probability of saving a life, the expected number of life years saved, the expected future wages saved—focus on the effect of life extension on consumption or income. A utility theory approach considers both aspects, but is difficult to use. In "Discounting Risk Reduction Benefits With Discounted Longevity," Robert Bordley shows how to develop an easy to use utility-based measure, which decouples the consumption and income effects in a fairly simple way. The criterion, discounted years of life saved, represents a compromise between "lives saved" and "expected number of life years saved."

Solving Problems With Risk

An investor may want to allocate funds among several projects. Each project is expected to yield a reward, with the objective that the reward will achieve a given amount, called the *target.* This problem is relatively easy to solve when rewards are deterministic, but is hard in more realistic settings when the rewards are stochastic and the investor wants to maximize the probability of attaining the target. The deterministic version can be solved by dynamic programming. However, dynamic programming fails in the stochastic version when risk criteria are involved. When rewards are normally distributed, it may be formulated as a nonlinear integer problem. In "Risk Criteria in a Stochastic Knapsack Problem," Mordechai Henig shows that by combining dynamic programming with a search procedure this problem can be solved relatively quickly.

Refining Diffuse Approximations

Queues are used to model a variety of systems, such as computer and communications networks. Generally, it is impossible to obtain simple, exact expressions for measures of performance, such as the queue length distribution or the distribution of unfinished work. So, diffusion approximations have become a popular tool for analyzing queueing models. These approximations often yield simple results for untractable systems. Diffusion approximations are asymptotically correct in the heavy traffic limit. If we view the diffusion approximations as the leading term in an asymptotic expansion of a performance measure, it is natural to compute some corrections to this leading term. In "Refinements to Heavy Traffic Limit Theorems in Queueing Theory," Charles Knessl shows how to use singular perturbation methods to compute these correction terms. This leads to refinements that increase the numerical accuracy of diffusion approximations.

The Newsboy Faces Delay

Very often in situations of demand uncertainty, a production or supply decision has to be made before the most recent demand figures are available. The traditional newsboy does not learn from experience, and so, ordering in advance is no problem. However, if learning is possible, time delay is relevant (and costly). If demand is deterministic and fixed, the newsboy learns in each period whether there was a sellout or he finds the exact demand. In "A High-Low Search Algorithm for a Newsboy Problem With Delayed Information Feedback," Diane Reyniers discusses how the newsboy should set supply levels to minimize the worst case cost if there is a time delay.

Exorcising Dimensionality

River basins with multiple reservoirs and inventories at multiple locations are modeled as dynamic systems with vectors as states. As the dimension of the state increases, there is often a rapid growth in the computation and storage needed to find a solution. This relationship is the *course of dimensionality,* which can sometimes be *exorcised* with a *myopic* solution—analyzing a static optimization problem instead of a dynamic program. In "Myopic Solutions of Dynamic Models," Matthew Sobel enlarges the class of models that are known to have myopic solutions. General results are presented for dynamic programs and sequential games. The power of these results is illustrated with an aquaculture harvesting model and a dynamic oligopoly model with advertising decisions.

Controlling Service Systems

Traditional control models for service systems emphasize policies based on state information about, for example, the number of jobs or the workload in a system. In practice, the controller may have some advance knowledge—for some future jobs, arrival epochs and processing requirements may be known. The controller may use this information to reject an arriving job in favor of a future one, or reallocate jobs among servers, depending on the objective. In "Look-Ahead Policies for Admission to a Single Server Loss System," Wim Nawijn analyzes the optimal admission policy to a system without waiting room, in which the controller has knowledge of the first future job and its arrival epoch. The policy tells when to reject an arriving job in favor of a future job in order to maximize the throughput of the system.

Closed Networks of Queues

The theory of queues may be applied to many real-life situations. Examples range from the prosaic—customers queueing in supermarkets or at airline counters—through the technological and esoteric—packets of information queued within satellites or data networks. The last example is chosen advisedly, as many developments in queueing theory have been motivated by problems in telecommunications, and the theory of queues has been used to resolve a variety of issues relating to the design and management of communication networks. The theory of networks of queues is an intriguing subject with obvious applications to communication systems. Unfortunately, even if a network has a simple mathematical solution, the theory is often difficult to apply due to the large number of system states involved. In "The Evaluation of Normalizing Constants in Closed Queueing Networks," J. J. Gordon describes a complex variable technique that allows one to carry out sums over large numbers of network states.

Priority Pricing Problems

Incentive and information problems arise in the management of computer systems where jobs have to be priced and prioritized. Jobs have diverse values, time values and service requirements. To maximize the efficiency of a system the manager may choose high-valued jobs and process them in prioritized order. In practice, however, the manager does not have full information of the parameters necessary to implement the optimal rule. Instead, she has aggregated statistics, while job-specific values are known only to the user. Because each user is oriented toward self-interest rather than to overall system optimization, eliciting this information from users is difficult. In "Optimal Incentive-Compatible Priority Pricing for the M/M/1 Queue," H. Mendelson and S. Whang develop a priority pricing scheme that aligns the incentives of users possessing private information with the overall optimum. Given the proposed priority pricing scheme, individuals find it in their best interest to follow the system-optimal rule.

Existence of Forecast Horizons

It is a common business practice to use rolling horizon procedures to solve infinite or long horizon decision problems. The finite horizon problem is solved and the first decision is implemented. The same process is repeated every period. Forecast horizons are of interest in a rolling horizon environment because they provide an upper bound on the number of periods to guarantee optimality of the first decision. In "Existence of Forecast Horizons in Undiscounted Discrete Time Lot Size Models," S. Chand, S. Sethi and J. M. Proth provide a forecast horizon for the discrete-time undiscounted dynamic lot size model in which demand and cost parameters are constant for a few periods. Their results for forecast horizons in an undiscounted optimization problem are obtained for the first time.

Dispatching in Manufacturing

Cellular manufacturing, a blend of group technology and flexible manufacturing, has been gaining popularity in modern industries. A typical cellular manufacturing system consists of several independent manufacturing cells and may be characterized as a stochastic parallel processing system. Determining the optimal dispatching schedules—the job orders dispatched to each cell—for such a system is a new and difficult problem, which has not adequately been analyzed. In "Optimal Dispatching in a Periodic Review Cellular Manufacturing System," John Liu develops a dynamic programming model to determine the optimal dispatching policy that minimizes the expected operational costs.

Positioning Stock in a Warehouse

A warehouse often consists of automated storage and retrieval devices in addition to conventional devices, such as shelving. A warehouse manager must decide which items to store in which devices and in what quantities. At first, the items most frequently requested would seem to be best suited to automated handling, but if they are requested in large quantities, or are larger in size, then excessive space may be

(continued on page 931)

IN THIS ISSUE (*continued from page* 746)

required. To minimize restocking and retrieval costs, it may be better to assign selected slower-moving but smaller items to automated equipment. This stock positioning problem motivated the more abstact problem considered in "Near-Optimal Solution of Generalized Resource Allocation Problems With Large Capacities" by Steven Hackman and Loren Platzman. They allocate resources to activities where each variable may be restricted, the objective function is additively-separable, and each activity uses at most one resource. They present a simple and efficient algorithm to solve a nonsmooth convex relaxation of this problem, and they prove that the solution will be near optimal when the number of activities is large.

Mixed Integer Linear Bilevel Programming

Decision making in large, hierarchical organizations rarely proceeds from a single point of view. Planning is an interactive process in which a central unit must coordinate the activities of lower level units to optimize system performance. The situation gets even more complicated when control is decentralized and all units act with some degree of independence. The resultant problem may be modeled as a sequential two-person game, known as a bilevel program. In "The Mixed Integer Linear Bilevel Programming Problem," Jim Moore and Jonathan Bard develop a series of algorithms for solving the case where each player controls discrete as well as continuous variables. They show that in most circumstances the standard branch-and-bound fathoming rules cannot be used, implying a need for extensive enumeration.

Bounds on Linear Programming

In linear programming problems of extreme size, or when the cost of obtaining individual columns of the coefficient matrix is restrictive, it may be acceptable to stop prior to reaching optimality. A bound on the optimal solution is desirable in such circumstances. In "A Note on Bounding a Class of Linear Programming Problems, Including Cutting Stock Problems," Alan Farley presents such a bound for a useful class of linear programming problems.

Are Schedules Changed by Buffers?

In production lines, finite buffers may be placed between consecutive machines to improve system performance. Is the optimal schedule changed by installing these buffers? In "Stochastic Minimization of the Makespan in Flow Shops With Identical Machines and Buffers of Arbitrary Size," M. Kijima, N. Makimoto and H. Shirakawa show that, under some processing time distribution assumptions, the forms of the optimal schedule in a flow shop are independent of buffer size.