Operations Research

In This Issue

To cite this article:
(1993) In This Issue. Operations Research 41(6):1010-1012. https://doi.org/10.1287/opre.41.6.1010

IN THIS ISSUE

Solving a Distribution Problem

Baxter Health Care Corporation worked with the authors of this *OR Practice* paper in explaining their problem, and in providing data and feedback on results. Baxter had developed a mixed integer programming model of their distribution system which involved shipping product groups from plants to intermediate facilities to distribution centers. The system contained 30 product groups, 7 intermediate facilities, 49 distribution centers, and 3 modes of transportation. In addition, some single sourcing and weight and volume restrictions further complicated the problem. Baxter's attempt to solve this problem using their formulation failed. It ran for more than 100 hours of CPU time on a mainframe computer without yielding an integer solution. Cynthia Barnhart, Ellis Johnson, George Nemhauser, Gabriele Sigismondi, and Pamela Vance approached the problem by using some integer programming reformulations, such as disaggregation, and a new way to model the weight-volume restrictions that reduced the number of variables. The new model had 15,128 columns, 27,441 rows, and 1,122 0-1 variables. It was solved in about 13 minutes on a workstation.

R&D Project Portfolio Management

One reason for the crisis in industrial competitiveness in the USA is the low rate at which organizations successfully commercialize R&D projects. Corporate R&D project portfolios across the USA have a 14% successful commercialization rate nationwide. This *OR Practice* paper describes how the Gas Research Institute achieved a 28% successful commercialization rate over a 14-year period with the help of OR. William Burnett, Barry Silverman, and Dominic Monetta describe a multicriteria, probabilistic scoring model and associated methodology, and how that approach helped to support their many advisory boards' judgments and portfolio decisions. The contributions of this case study to OR are two-fold: few other OR systems (no others in the R&D project appraisal field) have as long a history of sustained use in a real-world organization, and few other OR systems have lead to as many economic, societal, and qualitative benefits. For these reasons, how the Gas Research Institute balanced the use of OR with other methods is a valuable case to document. The lessons learned should be instructive to other companies struggling to increase their competitiveness.

Modeling Multicriteria Decision Problems

In real-time decision making practice, decision alternatives are either too numerous to consider or not all immediately available. Moreover, decision makers are not completely sure of their true preferences under such circumstances. Application of traditional decision analysis in such an environment may not be especially useful and may be dysfunctional. P. Korhonen, H. Moskowitz, P. Salminen, and J. Wallenius propose a general algorithm and supporting theory for progressively modeling and solving multiple criteria decision problems in such situations, where decision alternatives can be introduced dynamically and knowledge of one's preferences is incomplete. The algorithm is based on progressively sampling the decision space, obtaining preference information from the decision maker, determining the chances of finding possibly or surely better alternatives, and based on this information, continuing the search or terminating it by making a final choice. The authors describe a computerized implementation and testing of the algorithm, as well as experiences in field testing and practice.

Scheduling a Flexible-Duration Rest Period

All previous modeling of work performance reported in the literature has focused on individual employees. In contrast, "Optimal Scheduling of a Flexible-Duration Rest Period for a Work Group" considers the situation where the work break must be taken simultaneously by all employees in a work group. Stephen Bechtold and Gary Thompson develop and evaluate a simple and efficient optimal solution procedure (EOSP) for this problem. EOSP generates optimal solutions to problems containing more than 2,000 employees in under two seconds on an 80386DX/20-based microcomputer.

Waiter, There is a Matroid in My Pizza!

The delivery man problem (DMP) is a variant of the traveling salesman problem (TSP): In the DMP, the objective is to minimize the average arrival time at various customer locations. Pizza delivery is a common example. Another application is the routing of automated guided vehicles through production cells in a flexible manufacturing system. A well-known machine scheduling problem can also be viewed as a DMP: The problem is to minimize mean flow time for jobs with changeover times, arriving at time zero, and to be processed sequentially on a single machine. The DMP is NP-hard in the strong sense and is more difficult to solve than the TSP. Matteo Fischetti, Gilbert Laporte, and Silvano Martello develop new theoretical results on the matroidal structure of a related class of combinatorial problems. These are used to obtain an enumerative algorithm for the DMP, a dominance criterion, and a heuristic algorithm. Computational experiments indicate that DMPs involving up to 60 customer locations can be solved to optimality, by far the largest size ever attained.

Improving Throughput in a Factory

As a part of the concern with time-based management in contemporary organizations, one technique often used to accelerate work in batch systems is transfer batches. Whereas the traditional approach transfers work through a system in a single batch, it is often effective to split a batch and to move the work in sublots. The optimal sizing of transfer batches is a complicated allocation problem, and equal-sized transfer batches are usually suboptimal. In "Basic Techniques for Lot Streaming," Dan Trietsch and Ken Baker provide a framework for this optimization problem and show how to find optimal solutions for a number of important cases with relatively simple procedures. These procedures not only provide insight into the solution of lot streaming problems, but they can be used to enhance existing production control software.

Routing Automated Guided Vehicles

When AGVs are used in a bi-directional network layout, they need to be routed in a manner that results in no collisions or dead-lock situations. Yet, they need to be routed in an effective manner from a production or manufacturing perspective. This poses a challenging traffic-control problem. Nirup Krishnamurthy, Rajan Batta, and Mark Karwan develop a column generation-based heuristic to solve the static version of this problem. Empirical results indicate that the procedure usually generates solutions that are within a few percent of a proposed bound, within reasonable computer time. The method provides an effective means of achieving real-time control of AGVs in a dynamic jobshop environment.

Improving Flow Performance by Scheduling

The production of silicon integrated circuits for microprocessors being constructed by a large telecommunications company involves complex queueing networks in the processing of lots in the production facility. In particular, multipurpose workstations, or hubs, processing different tasks for the same lot at various points in the production process, can create serious scheduling problems. Improved understanding of hub networks is crucial to improving the throughput performance of such manufacturing systems. In "Batching and Scheduling in FMS Hubs: Flow Time Considerations," Jeff Sidney and Dan Lane model the hub structure and present some results for dual measures of flow time: total flow time required for a fixed number of lots, and average flow time for all lots. Their results form the basis for further studies that enhance the understanding of the behavior of hub networks. Better understanding of hub networks will lead to improved performance of scheduling complex production networks.

Dynamic Scheduling of a Fluid Network

A central issue in the management and control of many engineering systems, e.g., manufacturing systems, is to allocate resources, machines and operators, for instance, among jobs that need to be processed by the resources. When the system scale is large, the discrete nature of the jobs can be ignored effectively, and the system can be treated as a network of servers (resources) that process continuous job flows. Different types of jobs can be viewed as fluid with different colors. Scheduling then takes the form of dynamically allocating the processing capacity of each server among different colors of fluid. Hong Chen and David Yao develop a systematic approach to solving such dynamic resource allocation problems, and demonstrate its relevance to scheduling in stochastic networks with discrete jobs, which is an intractable problem in general.

A Case for Time-Dependent Decisions

Controlled Markov chains serve as key models in many application areas of OR. However, existing algorithms for constrained optimization, which are based on stationary policies, are not feasible computationally for large state spaces, and cannot be used in adaptive applications. Eitan Altman and Adam Shwartz propose a class of nonstationary policies that allows the explicit computation of optimal policies for constrained Markov decision processes with average cost criteria. These policies are applied to problems arising in telecommunication networks, and to general constrained adaptive problems.

Maintenance for Systems With Identical Components

The problem of how to perform maintenance and replacement in systems with many identical components arises in various contexts, like the replacement of personal computers in a department, the maintenance of segments in a railroad track or asphalted highway, or maintaining the piers of a bridge. Optimal policies based on the complete age configuration of all individual components are not well structured nor easy to implement. On the other hand, group-wide block-replacement policies disregard the stage of the individual components completely. Frank Van der Duyn Schouten and Stephan Vanneste investigate two simple group-replacement policies which recognize the advantages and disadvantages of complete individual component information. The proposed policies use detailed and global component information at different levels, and they compromise between the two types of policies mentioned above.

Control Congestion in a Queue

Workstations are often arranged along the path of a conveyor in an ordered sequence in a manufacturing system. Jobs carried by the conveyor pass the station and are withdrawn by the first station that is available. If a job cannot find a free station it is placed in a storage area until a workstation becomes available. Such an operations protocol also arises in other contexts, such as data communication networks and transportation systems. The issue of congestion control through admission under these circumstances has not been addressed adequately. Susan Xu and J. George Shanthikumar use a new approach to find the optimal admission policy that maximizes the system's profit. Additionally, their approach has applicability to other Markov decision problems besides admission control.

Global Conversion Results

The problem of locating a new facility relative to a set of customers to minimize transportation costs has been studied extensively. In the majority of cases, it is assumed that travel distances are either Euclidean (straight line) or rectangular (city block), even though Love and Morris proved (1972, 1979) that the l_p norm provides a more accurate model of travel distances. The Weiszfeld iterative procedure was proven to converge by Kuhn (1973) for Euclidean distances. However, whether or not the generalized Weiszfeld procedure converges for l_p distances remained an open question. J. Brimberg and R. Love achieve these results: The generalized Weiszfeld procedure is guaranteed to converge when p takes on any value in the closed interval from 1 to 2, covering a continuous spectrum of distance measures from rectangular to Euclidean, and convergence is not guaranteed when p is greater than 2.

Convergence of the HAP

The Euclidean single facility location (ESFL) and the Euclidean multifacility location (EMFL) problems are widely studied, and have applications in transportation and logistics. For the ESFL problem, Weiszfeld (1937) proposed a simple, closed-form iterative algorithm. Later, it was proved to be a convergent descent algorithm. In 1973, Eyster et al. extended Weiszfeld's idea and proposed a hyperboloid approximation procedure (HAP) for solving the EMFL problem. They conjectured that the HAP always converges. Since then, many papers have been published trying to prove the convergence of the HAP. In this paper, J. B. Rosen and G. L. Xue finally prove this conjecture.

Solving Minimax Problems Directly

In many practical contexts, such as facility location or assembly line balancing, one is often faced with the problem of minimizing the maximum value of some performance measure. The decision variables for these problems are typically logical, and the performance measures and constraints are often expressed as linear functions of the variables. Because of the minimax nature of these problems, they are normally solved by mixed integer programming. M. Diaby presents a method which exploits the pure-integer nature of these problems to solve them directly. The computational requirements of the method are several orders of magnitude smaller than those of the traditional mixed integer programming approach.