Operations Research

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

In This Issue

To cite this article:
(1994) In This Issue. Operations Research 42(1):2-193. https://doi.org/10.1287/opre.42.1.2

IN THIS ISSUE

The Age of Optimization

During the past decade new advances in algorithms have been as important as the impressive advances in computer technology. Using the new interior-point algorithms and advanced implementations of simplex methods we can now solve linear programs with more than a million variables and thousands of constraints. Preprocessing and polyhedral theory have yielded at least an order of magnitude improvement in branch-and-bound algorithms for solving mixed integer programs. Moreover, these algorithmic advances have been incorporated in commercially inexpensive software that is readily available, easily portable, and supported by a variety of systems that make it possible for unsophisticated users to input and check their models and obtain understandable outputs. In this *OR Forum* paper, which is based on his Morse Lecture given at the May 1993 TIMS/ORSA meeting in Chicago, George Nemhauser reviews the modern history of optimization, surveys some recent developments (illustrating them with an application in the airline industry), and offers a view of the future.

Reducing Costs by Millions of Dollars

A new discounting approach called business volume discounting is emerging. A. Sadrian and Y. Yoon describe a procurement decision support system (PDSS), which is based on this new approach. Using PDSS, several telephone companies have reduced their procurement costs by millions of dollars. In business volume discounting, the discount is based on the total dollar amount of the business awarded to the supplier for purchase of various products made from that supplier. This is a dramatic departure from traditional quantity discounting, where the discount is based on the quantity of the product purchased, and there is a separate discount schedule for every product. Business volume discounting is gaining momentum because it consolidates the buyer's purchasing power and increases negotiation leverage with the supplier; it gives the buyer a direct incentive to increase the seller's total revenue; it reduces the cost of managing suppliers by reducing the supplier base; and it makes more sense when the setup cost associated with the manufacturing of a product (or providing a service) is negligible because the setup cost drives quantity discounts.

Real-Time Data Processing

Optimal policies that maximize the availability of the system under uncertainty traditionally are in the center of interests of the OR community. In "Optimal Real-Time Date Acquisition and Processing by a Multiserver Stand-By System" Joseph Kreimer and Abraham Mehrez propose an optimal operation policy which does not depend on the arrival pattern of processed data. Moreover, it maximizes the system's availability during any finite period of time. The results of their research were applied to improve the reconnaissance activity of remotely piloted vehicles (RPV). Their approach may be of use in many areas, such as production control, signal processing, traffic surveillance, and high-performance data acquisition and processing.

Ten Years of Practice

At this transition period, when Mike Rothkopf has resigned as the area editor of *OR Practice* to become the editor of *Interfaces*, and Margaret Brandeau has assumed the area editorship of the practice section, it is appropriate to provide our readers with a summary of the papers that have appeared in the section since its introduction as a separate feature ten years (volumes) ago. As a result, the special note that appears behind the practice articles, was prepared at the request of the Editorial Office by Mike Rothkopf, who has served as area editor since the inception of the section. He examines and interprets the nature of the papers that the section has attracted and argues that they suggest vigor in the practice of operations research.

Library Applications

We are living in an information age. The classical notion of a library is very much part of it. It will serve as a worthwhile contributor to enhancing our quality of life and of knowledge for the foreseeable future. Yet, libraries compete for the dwindling resources in the public and the private sectors. If done right, OR/MS can help library management to better cope with the opposing goals of service enhancement and cost containment. In their survey paper, Arnold Reisman and Xiaomei Xu show that there is much to be done. The payoffs to society are great. Yet, the recent record of OR/MS accomplishment in this area

is especially sparse. This then presents a challenge to the OR/MS community.

Optimal Bus Scheduling

The planning process of a transportation system is usually divided in two phases: trip definition and vehicle scheduling. Trip definition is based on demand and marketing aspects and deals with local considerations. Vehicle scheduling considers the coordination of the whole system and the network of routes, trips and vehicles. The scheduling problem becomes more difficult when the vehicles are provided by several depots. C. C. Ribeiro and F. Soumis give a new formulation for the multiple-depot vehicle scheduling problem with side constraints, whose continuous relaxation is amenable to be solved by column generation, which provides a very tight lower bound. This approach makes it possible to exactly solve to optimality problems with up to 300 trips, opening the possibility of solving to optimality many other real-world problems in transportation and network planning.

A New Twist on Cutting Planes

Cutting plane methods for solving integer programs have generated considerable interest in recent years due to remarkable computational results on a wide variety of difficult problems. In "Fenchel Cutting Planes for Integer Programs," E. A. Boyd proposes a cutting plane method more closely related to Lagrangian relaxation than to more common facet generation techniques. Similarities and differences with alternative cutting plane methods are discussed, and computational results that demonstrate the applicability of the cutting planes are presented.

Fast Network Optimization

Recent advances in computer designs present new opportunities for expediting the solution of important problems. Richard Barr and Betty Hickman report on a new parallel implementation of the primal simplex method for minimum cost network flow problems that decomposes both the pivoting and pricing operations. The self-scheduling approach is close in speed to the best serial code when using one processor, and is capable of substantial speedups as parallel computing units are added. An in-depth computational study of transportation and transshipment problems verified the effectiveness of this approach, with results on a 20-processor 80386-based system that are competitive with, and occasionally superior to, massively parallel implementations using tens of thousands of processors. A microanalysis of the code's behavior revealed unexpected sources of the occasionally superlinear speedup.

Pick-and-Place Machines

Printed circuit boards (PCBs) are an important module of electronic parts, and consequently, concurrent pick-and-place machines have been designed to produce the required PCBs at the desired volume, quality, and cost. An important problem is the staging of these machines. Suboptimal solutions of the staging problem can severely degrade the throughput rate of the pick-and-place machines. R. Ahmadi and P. Kouvelis have formalized the staging problem as an integer programming problem and have optimally solved the single product staging problem using an effective branch-and-bound procedure. The multiproduct staging problem was solved using an algorithm which capitalizes on the computational efficiency of the single product case. The computational experiments for moderate size problems suggest the satisfactory quality of the proposed solution technique.

Scheduling Raw Materials

Economic lot scheduling concerns scheduling the production of several items on a single facility to meet constant external demands. Switching production of these items requires downtimes and costs for the setup. The focus is on average, long-run operating costs in the presence of setup time constraints. Research on this topic universally concentrates on the single production facility, separated from its surrounding processes. In practice, raw materials may constitute a significant part of the costs of production, and schedules that take these into account can be markedly different from those that neglect raw materials. Dev Joneja and Guillermo Gallego consider both the facility and the raw material. The authors formulate a planning model which provides a sharp lower bound on the cost of the schedule, and develop a plan for the finished items and the raw materials which is close to optimal. They also address the issue of obtaining feasible schedules that perform well in practice and are proved to be close to optimal in the worst case.

Scheduling: Consider Starting (Too) Early

The just-in-time concept in manufacturing has triggered a lot of machine scheduling research on balancing early against tardy completions of jobs. Tight due dates complicate this balancing. Allowing the machine to start early removes the complication, but leads to infeasible schedules. Han Hoogeveen, Henk

Oosterhout and Steef van de Velde put this relaxation in a mathematical context for the case of a tight due date that is common to all jobs. This gives rise to the so-called Lagrangian dual problem, which is identical to the original problem in all aspects but one: The machine may start early at an extra cost. The solution to this has two nice properties. First, it is easy to make it feasible at little extra cost; this property gives rise to an approximation algorithm with an excellent empirical and worst-case performance. Second, its solution value is a strong lower bound on the optimal solution value of the original problem. These two properties together lead to an effective optimization algorithm for the original problem.

How Machine Unreliability Affects Performance

The fundamental problem in production-inventory (PI) systems is to relate the production process to the demand process to ensure that while shortages are kept at a low desirable level no excessive inventory is built up. PI models in the literature vary with respect to characteristics of the production process: the demand pattern, customer response to shortages, and inventory review procedures. But one frequently neglected element is the imperfection of production systems. Production machines can fail, disrupting production plans, and resulting in insufficient inventory and consequent shortages. M. J. M. Posner, M. Berg and H. Zhao incorporate this reliability factor into the PI analysis. This is done for a core model whose relatively simplified assumptions cover the different processes involved: production, demand, and reliability, and then for a host of modeling extensions which are obtained by relaxing or altering these assumptions.

Prediction for Queueing Networks

Two popular techniques for the analysis of queueing networks are decomposition approximation and heavy traffic analysis. In heavy traffic analysis, queue lengths and workload levels are generally approximated by a multidimensional reflected Brownian motion (RBM). Performance analysis of the network is then reduced to finding the stationary distribution of the RBM. Decomposition approximations, on the other hand, typically analyze stations in the network separately, treating each as a single queue. J. Dai, V. Nguyen and M. I. Reiman present a hybrid method for analyzing generalized Jackson networks that employs both decomposition approximation and heavy traffic analysis: Stations in the network are partitioned into groups of bottleneck subnetworks that may have more than one station; the subnetworks then are analyzed sequentially with heavy traffic theory. This approach avoids issues of computational complexity associated with the method of heavy traffic analysis while yielding approximations that are more sensitive than those of traditional decomposition approximations. Their results suggest that this hybrid method generally performs better than other approximation techniques.

Sensitivity Analysis for Reliability Systems

The designer of a reliability system may want to know the sensitivities of the system with respect to various parameters, such as component failure and repair rates. A designer may validate the model and identify critical components by using sensitivities. For example, a designer with a limited budget may want to improve the components corresponding to the largest sensitivities to gain the greatest improvement in overall system performance. Analytic methods are not practical for solving the mathematical models that arise in this setting, so simulation is often used. However, it is difficult to obtain accurate estimates of the system performance and its sensitivities using standard simulation methods. Therefore, the simulator must employ variance reduction techniques to improve the estimates. Marvin Nakayama, Ambuj Goyal and Peter Glynn give an accurate, low-cost method of estimating sensitivities during a simulation. Their method, which employs likelihood ratio gradient estimators in conjunction with importance sampling, has been implemented in SAVE, an IBM-internal software package, and used to model and analyze numerous large-scale reliability systems.

Analyzing the $\Sigma PH_i/Ph/1$ Queue

Queueing network models are useful for analyzing the performance of jobshops like manufacturing systems. Unfortunately, exact analysis can be carried out only for a limited class of networks. Approximations are needed for general networks. An approximation procedure that has been effective for estimating the first moments of queue lengths is the decomposition approach. To estimate higher moments of the performance measures, it is necessary to refine the decomposition methods. G. R. Bitran and S. Dasu analyze queues to which several classes of customers arrive and determine the performance measures for each class. The developments in this paper are useful for estimating higher moments in networks and to refine several approximations that are carried out within the decomposition framework.

(*continued on p. 193*)

(In This Issue continued from p. 4)

Mean-Variance Tradeoffs

Manufacturing, telecommunications, and health care are among the industries in which dynamic programming models are applied. In spite of empirical evidence that people are sensitive to risk, most theory and algorithms assume risk neutrality. However, a small recent literature employs risk-sensitive optimization criteria for MDPs (Markov decision processes). MDPs include most models to which dynamic programming is applied. In "Mean-Variance Tradeoffs in an Undiscounted MDP" Matthew Sobel introduces efficient algorithms to explore tradeoffs between the long-run average reward per time period and the long-run variance. Most decision makers prefer high average rewards and low variances. Sobel shows that mean-variance tradeoffs can be explored with essentially the same computational effort as is needed to optimize the average reward.

Mean-Variance Tradeoffs: The Unichain Case

For a major portion of the work in infinite horizon Markov decision processes, standard objective function criteria have been the expected limiting average reward per unit time, or the expected discounted total reward (averaged or discounted) may yield optimal policies which are unacceptable to a risk-sensitive decision maker. Such a decision maker may well prefer to use a criterion which incorporates risk in the decision process. The model developed by Ken-Jun Chung incorporates risk in the decision process. Both mean and variance are simultaneously considered as criteria to select the Pareto optimal policy. The analysis provides insights into mean-variance tradeoffs and discusses the relation between Pareto optimal policies and nonrandomized policies.

Open Shop Problems

Open shop problems occur in the field of production, for instance, planning repair and maintenance work. Most of the open shop problems belong to the NP-hard class. Polynomial algorithms have been developed for such problems with unit processing times. By means of a new model we can solve open shop problems with unit processing times, where a deadline is imposed for every job. Thomas Tautenhahn gives a fast algorithm for minimizing total completion time and improved algorithms for computing the minimal number of late jobs and maximum lateness.

On the Busy Period of an $M/G/1/K$ Queue

This note concerns the length and the number of customers served during the busy period in an $M/G/1$ queue with a finite capacity. The busy period is a fundamental quantity in a queue that can be used for analyzing the performance of many complicated queueing systems in practical applications, especially in telecommunication networks. H. Takagi and R. O. LaMaire derive a new expression for the joint distribution of the length of a busy period and the number of messages served during that busy period in an $M/G/1/K$ queue. Second, they relate the result with Harris' result, which appeared in a previous issue of *Operations Research*, and correct his error.