Operations Research

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

In This Issue

To cite this article:
(1992) In This Issue. Operations Research 40(1-supplement-1):S2-S177. https://doi.org/10.1287/opre.40.S1.2

IN THIS ISSUE

Comment from the Editor

This supplemental issue is being published to reduce the backlog of accepted articles, and therefore decrease the time between acceptance and the publication of new articles. The issue is devoted entirely to papers that have been accepted in the **Optimization Area** (Jan Karel Lenstra, area editor). This is an opportunity for the readers of *Operations Research* to see and evaluate a large sample from the papers that we are publishing in this area. We invite comment on the topics covered and the quality of the manuscripts in this issue.

H. DONALD RATLIFF, Editor
Atlanta, Ga., December 1991

An Improved Network Simplex Algorithm

The minimum cost network flow problem has been used to model a variety of settings including, manufacturing, communications, transportation and distribution, military science, and computer science. An excellent approach for the problem in practice is the network simplex algorithm: The entering and leaving variables are selected according to the simplex algorithm of linear programming, and the performance is dramatically sped up by calculating values using network and tree data structures. Although the algorithm is excellent in practice, its best implementations in practice are terrible in the worst case. They run in exponential time even if restricted to easy special cases, such as assignment and shortest path problems. In "The Scaling Network Simplex Algorithm," R. K. Ahuja and J. B. Orlin give an easily implementable, efficient rule for the algorithm that is similar to the best rules used in practice, but which has enhanced performance guarantees.

Habitat Dispersion in Forest Planning

The problem of detailed placement in forest management has received a good deal of attention recently. There are reasons like scenic beauty and wildlife protection that force no cutting in large areas for periods of approximately 10 years. Up to now, this type of constraint had been incorporated in optimization models and solved with different heuristics. In "Habitat Dispersion and the Stable Set Problem," Francisco Barahona, Andrés Weintraub and Rafael Epstein tackle the forest planning problem with a column generation approach. Each new column is obtained by solving a maximum stable set problem in a graph associated with the areas to be cut. The problems that have to be solved in each iteration are treated with a cutting plane approach based on a partial characterization of the stable set polytope. This is a relatively easy task because the graphs that arise are not very complicated. Practical experience is presented with forests in southern Chile.

Bounding the Quadratic Assignment Problem

The quadratic assignment problem (QAP) is difficult. A lot of real-life problems can be formulated as a QAP, such as plant location, schedule synchronization and keyboard design. The most effective solution techniques are the branch-and-bound approaches, which fail to solve problems of a size larger than 15. The main reason for this failure is the lack of good lower bounds. In "A New Lower Bound for the Quadratic Assignment Problem," P. Carraresi and F. Malucelli propose a new lower bound. A sequence of equivalent reformulations of the problem is also provided, which implies that other lower bounds can be applied in conjunction with the proposed one, yielding better results. Promising experimental results are presented.

On the Classroom Assignment Problem

Most people know that assigning classrooms is a tough problem in most large universities. However, it is somewhat of a surprise to discover what really makes it difficult. In particular, the fussy teacher who wants special facilities and particular seating styles makes the usual space problem a lot harder. Michael Carter and Craig Tovey have gone to some trouble to make their paper, "When Is The Classroom Assignment Problem Hard?" readable and interesting for the practitioner as well as the theoretician. The style makes it accessible to people who are actually solving room assignment problems and no background in complexity theory is necessary.

A New Set Partitioning Approach That Saved Millions of Dollars

Although many scheduling problems can be formulated as set partitioning problems (SPP), most real-life models are too large to be solved optimally. One example is the crew pairing problem (the assignment of airline crews to flights such that each crew starts and ends at the same location) in the airline industry. This is an enormously complex problem because the variables are innumerable. At American Airlines (AA), the approach used to solve this problem is to decompose it into small SPP subproblems and solve thousands of them iteratively. The quality of the solution improves as the subproblem size increases or the total number of subproblems solved increases. In fact, a 1% improvement in AA's crew utilization translates into $13 million annual savings. Applying the techniques described in "A Multiplier Adjustment Approach for the Set Partitioning Problem" (by Thomas Chan and Candace Yano), AA was able to increase the iteration speed and solve larger subproblems. Together with other innovative strategies, AA has been able to save over $20 million annually through better crew pairing solutions. However, there is still plenty of room for improvement.

Transmitting Flows Safely

In many practical situations, there are a few kinds of distinct commodities flowing simultaneously in a transport network. Each commodity has its own source and sink. Previous researchers resolved the problem by sending these commodities from the sources to the sinks with minimum transmission cost. The solution indicates the amount of flow that should be sent through a path. In "Multicommodity Network Flows With Safety Considerations," the difficulty of a path is considered. Semantically, this difficulty can be interpreted as the worst road condition in that path. The authors assume that each commodity has a tolerance capability, which is the worst road condition the commodity can tolerate without harm. A path is unsafe if the degree of path difficulty is harder than the tolerance capability of the commodity. Otherwise, the path is safe for this commodity. This research has two goals: to send these commodities in the cheapest manner, and in the safest way. Y. Chen and Y. Chin propose a model which achieves both goals simultaneously, and show that the model is equivalent to the conventional model.

Locating Facilities That Interact

Multiple facilities, which interact with existing demand points as well as with each other, are to be located on a transport network. In addition to the interaction costs, there are fixed location costs. The objective is to minimize the total cost or to minimize the maximum of the cost of each interaction and the fixed location costs. These are important location problems with applications to manufacturing, distribution, and telecommunications. Past research has dealt with exploiting the structure of the transport network. The problem can be solved efficiently when the network is a tree. In "m-Median and m-Center Problems With Mutual Communication: Solvable Special Cases," Dilip Chhajed and Timothy Lowe focus on the structure of the interactions between pairs of facilities. This structure is captured via a flow graph. They give an algorithm to efficiently solve the problem when the flow graph has a special structure. Thus, this paper provides a different approach to identifying and exploiting special structures in location problems.

Scheduling to Minimize Expected Makespan

Scheduling jobs on parallel machines under precedence constraints having a forest structure is a fundamental problem of stochastic scheduling theory. Job running times are not known in advance, but they are known to be independent samples from a given exponential distribution. The objective is to schedule the jobs to minimize the expected value of the latest finishing time, i.e., the makespan. Much has been written about the case where precedence constraints form an in-forest, in which a job cannot be run until all of its descendants have finished. In "On the Optimal Stochastic Scheduling of Out-Forests," Edward Coffman and Zhen Liu study the converse problem of out-forests, in which a job must finish before any of its descendants start. They exhibit a broad subclass of out-forests for which a simple greedy policy is optimal. Their results, along with several counterexamples for the general case, give deeper insights into this difficult problem.

Scheduling Two Types of Processors

Scheduling problems that involve resources of different capacities, such as ships or classrooms, arise frequently in practice. An essential feature of this type of problem is that a processor can be assigned only to a subset of the jobs, and a job can be done only by a subset of the available processors. Many of these problems are NP-complete and there are no efficient algorithms to solve them. In "Fixed Job Scheduling With Two Types of Processors," Reddy Dondeti and Ham Emmons examine a scheduling problem involving two types of processors and three types of jobs with fixed start and end times. A polynomial algorithm is presented for finding the minimal cost combination of the two types of processors required to complete the jobs.

Trauma Care Evaluation

The leading cause of death in the United States for people between the ages of 1 and 38 is accidental injury. Approximately 150,000 of the 1 million people who are severely injured each year do not survive. Moreover, an estimated 15% to 30% of these deaths are the result of delayed or inadequate trauma care. Therefore, it is important for individual medical institutions to be able to evaluate the quality of their services for trauma patients. One measure of the effectiveness of institutional trauma management is based on collected patient data and requires the computation of a standard normal statistic. However, a potential weakness of the measure arises from incomplete patient data. In "Bounds on a Trauma Outcome Function Via Optimization," J. Falk, S. Palocsay, W. Sacco, W. Copes and H. Champion describe a method for determining upper and lower bounds on the value of the standard normal statistic leading to the optimization of a nonconvex ratio function. They apply methods of fractional programming and global optimization to efficiently calculate the desired bounds. The resulting solution method is being integrated into the trauma institution evaluation procedure.

Scheduling Fixed Jobs

In a consulting project undertaken for the town of Ravenna, Italy, Matteo Fischetti, Silvano Martello and Paolo Toth implemented a system for the classical bus driver scheduling problem: Given the transit company's schedule for a particular day, determine the minimum number of driver duties that meet the schedule and satisfy requirement imposed by the union contract and company regulations. The constraints fall into two categories: spread-time constraints that limit time from the beginning to the end of a driver's day, and working-time constraints that limit the total number of hours in a shift. The requirements differ because drivers typically work split shifts. Solving the problem without either requirement was easy, while adding one produced a difficult problem with interesting theoretical properties. This is presented in the paper, "Approximation Algorithms for Fixed Job Schedule Problems."

Linearization for Computational Efficiency

Economical linearization of nonlinear combinatorial optimization problems is crucial in handling many large-sized, real-life problems, such as facility layout, R&D planning, communication networks, and capital budgeting. Lack of effective linearization procedures usually leads to the use of specific heuristics developed for particular problems, and therefore, might not be effective in solving other types of problems. In "A Linearization Procedure for Quadratic and Cubic Mixed Integer Problems" Muhittin Oral and Ossama Kettani provide a linearization method which is considerably more economical and offers an opportunity to use commercially available codes more effectively. Researchers may find the method useful in developing algorithms to solve nonlinear combinatorial problems, whereas practitioners may reformulate decision problems having nonlinear integer expressions, so that a commercially available code might be used directly and effectively.

Breaking Ties in Optimization

A decision-maker is considering two possible immediate actions. To choose one, he considers their consequences after 1 year, 5 years, 10 years, and so on, hoping to see if one appears to be better in the long run. But no clear winner emerges. When may he stop the evaluation and simply pick one of the two? Can he do so in confidence that a look further into the future will not help to distinguish between the two actions? In "A Tie-Breaking Rule for Discrete Infinite Horizon Optimization," Sara M. Ryan, James C. Bean and Robert L. Smith develop a procedure to select an optimal beginning to a sequence of decisions without relying on the common assumption of a uniquely optimal initial decision. They illustrate the method with examples from production planning and capacity expansion, including one which, due to multiple optima, could not be solved by any previous algorithm.

Capacity Expansion of Chemical Processes

Problems of synthesis and capacity expansion often arise in communications networks, electric power generation systems, road and computer networks, and chemical processes. These problems are usually

Continued on p. S177

IN THIS ISSUE (*Continued from page S4*)

difficult to solve because they are combinatorial in nature and involve economies of scale in a dynamic economic environment. In "Reformulation of the Multiperiod MILP Model for Capacity Expansion of Chemical Processes," N. V. Sahinidis and I. E. Grossman describe nontrivial integer programming formulations for long-range planning in the process industries. Even though these formulations require more constraints and variables than is standard, they generally give rise to tighter LP relaxations, reducing the computational effort of branch-and-bound methods. Results for the synthesis and capacity expansion of chemical complexes indicate that the solution of real-world problems can be obtained within reasonable computational time by using the proposed formulations. Similar formulations may be beneficial for solving multiperiod capacity expansion problems in other dynamic networks.

A New Method for an Old Problem

The economic lot sizing problem was introduced more than 30 years ago in a seminal paper by H. Wagner and X. Whitin. Their approach to solve the problem is a classical application of dynamic programming. In the past decades, much research has been devoted to designing faster variants of the WW-algorithm. However, those improved methods did have a lower computational complexity than the original algorithm. Until recently, it was generally believed that such methods did not exist. Then, in 1989, several groups of researchers developed, independently and almost simultaneously, solution methods that have a lower complexity. In "Economic Lot Sizing: An $O(n \log n)$ Algorithm That Runs in Linear Time in the Wagner–Whitin Case," Albert Wagelmans, Stan van Hoesel and Antoon Kolen present one method. Their algorithm is particularly interesting because of its appealing geometric interpretation.

Choosing Bonds

Linear programming and other optimization models are widely used to help select bond portfolios. In just over a decade this has become an active, vibrant area of applications. Many institutional investors use three models routinely. Hundreds of professionals are employed to tend the models; consultants abound. New formulations are constantly being refined and marketed. "The Structure of Structured Bond Portfolios," by Paul Zipkin, provides an overview of the field, focusing on the structures of the models themselves. There is considerable and often confusing variety among the popular formulations; the paper sorts the basic similarities and essential differences. Also, the author sketches the context in which the models are used, suggesting why they appeal to investors, as well as some of their practical limitations.

One-Half Approximation Algorithms

The problem of partitioning or clustering the nodes of an edge weighted graph has been of interest to the OR community for many years. Applications in engineering, business, and computer science abound. Unfortunately, the problem is theoretically intractable, and in practice is proving to be troublesome to solve. In "One-Half Approximation Algorithms for the k-Partition Problem," Thomas Feo, Olivier Goldschmidt and Mallek Khellaf develop several simple heuristics for different values of k. They prove that these methods are guaranteed to efficiently provide solutions which are in the worst case one-half weight of an optimal solution. The objective is to maximize the sum of the weight of all edges in the constructed clusters. Interestingly, the equivalent partitioning objective of minimizing the sum of the weight of the edges not in the clusters may not possess an efficient approximation algorithm.