Book Selection

Linear Network Optimization: Algorithms and Codes
DIMITRI P. BERTSEKAS
MIT Press, Cambridge, Massachusetts, 1991. xi + 359 pp. £35.95
ISBN 0 262 02334 2

A large proportion of practical optimization problems may be formulated as linear network flow problems, or something similar. This, together with the fact that Dimitri Bertsekas has made important contributions to the area means that the book automatically deserves serious consideration. The reader soon discovers that this potential is realized via the author's authoritative and refreshingly clear style. Indeed, the author goes to such pains to tie down every detail that readers with some prior knowledge of the area might, at times, find themselves urging him to proceed more quickly.

The prerequisites required of the reader are fairly modest. Knowledge of standard simplex theory to the level of artificial variables, degeneracy and duality including the complementary slackness relations, and probably an acquaintance with the classical transportation problem, would provide a bare minimum.

Chapter 1 introduces graphs, flows, shortest path and assignment problems, and provides a foretaste of concepts used in the main part of the book. However, more than a minimal experience is required for a full appreciation of the ideas and methods presented.

The methods designed fall into three groups, primal cost improvement (Chapter 2); dual ascent (Chapter 3); auction algorithms (Chapter 4). The most successful primal cost improvement methods in practice are specialized versions of the simplex method. This derives from the fact that for the minimal cost flow problem (which includes various other network flow problems as special cases), a basic solution corresponds to a spanning tree, T, of the network. A simplex iteration involves:
(a) selection of a non-tree arc (i, j) (non-basic variable) which together with T forms a cycle C;
(b) pushing flow around C until flow in some arc $(k, 1)$ reaches zero or the capacity of $(k, 1)$;
(c) removing $(k, 1)$ to form a new spanning tree (basis).
The algorithm can work entirely with the network and no explicit tableaux are needed.

By contrast, dual ascent methods keep primal–dual solution pairs satisfying complementary slackness, and the dual solution (feasible) is improved until primal feasibility is attained. The group includes the classical 'primal–dual' method and the very efficient relaxation methods introduced by Bertsekas.

The third group, based around an approximate form of complementary slackness, called ε-complementary slackness, includes auction and ε-relaxation methods which are associated with Bertsekas. He claims that, although ε-relaxation is typically outperformed by the relaxation method, it is better suited for parallel computation than the other minimum cost flow methods described in the book. This interesting observation may be followed up in another book by the author and J. N. Tsitsiklis[1].

The last, rather short, chapter is devoted to performance comparisons between methods. Finally, there are some 90 pages of state-of-the-art Fortran codes for selected methods. (These may also be obtained on diskette from the author.)

The book has many worked examples, and exercises which provide interesting extensions to the theory presented. It is probably too specialized for most undergraduate courses but should be of great value to researchers and postgraduate students. I have gained much from reading this book and can thoroughly recommend it to others interested in the subject area.

T. B. BOFFEY

Reference

1. D. P. BERTSEKAS and J. N. TSITSIKLIS (1989) *Parallel and Distributed Computation*. Prentice-Hall International, New York.