

Approximating a 2-edge-connected subgraph problem

Higher-Confidence Integer Programming

Robert Carr and Ojas Parekh
Sandia National Laboratories
P.O. Box 5800, Albuquerque, NM 87185

Abstract

We consider the following network design problem. We are given a graph. We wish to select a minimum-cost subset of edges such that removing any edge leaves all pairs connected. This is a 2-edge-connected network design problem. Each edge has two costs, one for building a single connection and one for building two parallel independent connections. There is a buy-at-bulk discount for the parallel edges. That is, it costs less than twice the cost of the single channel. Our problem is a special-case of the buy-at-bulk network design with protection problem. Both the general problem and our variant are NP-complete, and thus it is reasonable to find provably-good approximation algorithms. An α approximation for an optimization problem is a provably efficient algorithm that guarantees a solution no more than α times optimal.

Antonakopoulos *et. al.* [1] give the best currently known approximation algorithm for the general buy-at-bulk with protection problem, which is actively used for designing optical communication networks. This general problem specifies a set of pairs of nodes to connect rather than requiring connection of all pairs. There is a bandwidth requirement for each pair. The network must be able to robustly and simultaneously serve all these pairs (e.g connectivity from New York to LA along with connectivity from New Orleans to St. Louis). Antonakopoulos *et. al.* give a (large, unspecified) constant approximation bound when all pairs have a common node and an $O(\log^3 n)$ approximation for the general case.

We consider a simplified buy-at-bulk with protection problem that requires connectivity of all pairs. But we assume each edge is capable of supporting all potential conversations among the n nodes. Thus this is an uncapacitated version. We give a $5/3$ -approximation algorithm for this case. The best previous bound for this problem was 2.

The algorithm is based on convex decomposition. To use this technique, we formulate the problem as an integer program (IP): a linear objective function subject to linear constraints, and variable integrality constraints. We drop the integrality constraints on the IP to get an efficiently solvable linear-programming (LP) relaxation. This satisfies the algebraic constraints, but selects edges fractionally. A convex decomposition takes a scaled version of the LP optimal solution, and expresses it as a convex linear combination of feasible integer solutions. Thus one of the feasible solutions approximates as well as the LP scaling factor. We find two independent decompositions, each scaling different types of variables non-uniformly. We then combine the two decompositions into a new decomposition with a stronger approximation bound. To our knowledge, this is the first time anyone has used this kind of multi-decomposition approach.

References

- [1] S. Antonakopoulos, C. Chekuri, B. Shepherd, and L. Zhang (Bell Labs), Buy-at-bulk network design with protection, *Mathematics of Operations Research*, Vol. 36, No. 1, February 2011, pp. 71–87.