A Polyhedral Approach to Network Connectivity Problems (Extended Abstract)

by
M. Grötschel (ZIB, Berlin)
C.L. Monma (Bellcore, Morristown)
M. Stoer (ZIB, Berlin)

Abstract

We present a polyhedral approach for the general problem of designing a minimum-cost network with specified connectivity requirements. This includes identifying classes of facet-defining inequalities and using them in a cutting plane approach for obtaining optimal or near-optimal solutions. Preliminary computational results with this approach are presented.

1 Introduction

This extended abstract focuses on the important practical and theoretical problem of designing a minimum-cost network with specified connectivity requirements. Our initial interest in this area was motivated by the problem of designing survivable "two-connected" topologies for fiber optic communication networks for the regional telephone companies; see [CMW89] for an overview. Work on the two-connected network design problem naturally leads to theoretical and algorithmic questions for network design problems with higher connectivity requirements. There has been a great deal of research activity in this area in recent years; for a survey of research in this area, see [GMS93] and [S91].

In this extended abstract, we describe polyhedral results, including natural integer programming formulations, classes of valid and facet-defining inequalities, and their associated separation problems for network design problems with higher connectivity requirements. We also report some preliminary computational results with a cutting plane algorithm on some random problems and on some real-world problems with higher connectivity requirements. This builds upon our earlier success with this approach for network design problems with low-connectivity requirements; see [GM90], [GMS89] and [GMS92].

In order to formalize the problem, we need to introduce the following notation. A set V of nodes is given which represent the locations that must be interconnected into a network. A collection E of edges is also specified that represent the possible pairs of nodes between which a direct link can be placed. We let $G = (V, E)$ be the (undirected) graph of possible direct link connections. Each edge $e \in E$ has a nonnegative **fixed cost** c_e of establishing the direct link connection. The graph G may have parallel edges but contains no loops. The cost of establishing a network consisting of a subset $F \subseteq E$ of edges is the sum of the costs of the individual links contained in F.

The goal is to build a minimum-cost network so that the required connectivity requirements are satisfied. To model these requirements, we introduce the concept of node types. For each node $s \in V$ a nonnegative integer r_s, called the **type** of s, is specified. We say that the network $N = (V, F)$ to be designed satisfies the **survivability conditions** if, for each pair $s, t \in V$ of distinct nodes, N contains at least $r_{st} := \min\{r_s, r_t\}$ disjoint $[s,t]$-paths. These paths may be required to be either edge disjoint or node disjoint depending on the application. We note that much of this work generalizes easily to the case where arbitrary connectivity requirements are given between every pair of nodes. However, we have found that node types capture the important aspects of practical problems while still providing an interesting theoretical and algorithmic framework. So we choose to present our results in this manner.

We introduce further symbols and conventions to denote these node- or edge-connectivity models: Let $r(W) := \max\{r_u : u \in W\}$; and let $\mathrm{con}(W) := \max\{r_{st} : s \in W, t \in V - W\}$ for any $W \subseteq V$. Let $G = (V, E)$ be a graph. For $Z \subseteq V$, let $\delta_G(Z)$ denote the set of edges with one endnode in Z and the other in $V \backslash Z$. It is customary to call $\delta_G(Z)$ a **cut**. For any subset of edges $F \subseteq E$, we let $x(F)$ stand for the sum $\sum_{e \in F} x_e$.

Consider the following integer linear program for a graph $G = (V, E)$ with edge costs c_e for all $e \in E$ and node types r_s for all $s \in V$:

(1.1) $\min \sum_{e\in E} c_e x_e$

subject to

(i)	$x(\delta(W)) \geq \text{con}(W)$	for all $W \subseteq V, \emptyset \neq W \neq V$;
(ii)	$x(\delta_{G-Z}(W)) \geq \text{con}(W) - \lvert Z\rvert$	for all pairs $s,t \in V$, $s \neq t$, and for all $Z \subseteq V\backslash\{s,t\}$ with $1 \leq \lvert Z\rvert \leq r_{st} - 1$, and for all $W \subseteq V\backslash Z$ with $s \in W, t \notin W$;
(iii)	$0 \leq x_e \leq 1$	for all $e \in E$;
(iv)	x_e integral	for all $e \in E$.

It follows from Menger's theorem that the feasible solutions of (1.1) are the incidence vectors of edge sets F such that $N = (V, F)$ satisfies all node connectivity requirements; i.e., (1.1) is an integer programming formulation of the node connectivity network design problem. Deleting inequalities (ii) from (1.1), we obtain, again from Menger's theorem, an integer programming formulation for the edge connectivity network design problem. The inequalities of type (i) will be called **cut inequalities**, and those of type (ii) will be called **node cut inequalities**. We call the associated problems the NCON and ECON problems, respectively.

2 Polyhedral Results

The goal of the polyhedral approach is to convert the integer programming problem (1.1) into a linear programming problem and solve it using LP techniques. This involves identifying classes of inequalities valid for the ECON and NCON problems, and algorithmically finding valid inequalities as necessary. For an overview of the polyhedral approach, see [P83], [GP85], and [PG85].

The integer programming formulation (1.1) provides natural classes of inequalities. For the cut inequalities (1.1)(i) and node cut inequalities (1.1)(ii), we can solve the associated separation problems in polynomical time and have (complicated) characterizations of when they are facets; see [GMS89] and [S91]. We will introduce further useful classes of valid inequalities.

The first such class is the class of partition inequalities that generalize the cut inequalities (1.1)(i). Let $[W_1 : \ldots : W_p]$ be the set of all edges having their endpoints in different sets W_i of the partition of V with $r(W_i) \geq 1$ for i = 1, ..., p, and let $I_1 := \{\, i \mid \mathrm{con}(W_i) = 1\,\}$ and $I_2 := \{\, i \mid \mathrm{con}(W_i) > 1\,\}$. Then the **partition inequality** induced by $\{W_1, \ldots, W_p\}$ is defined as

$$(2.1) \qquad x([W_1 : \ldots : W_p]) \geq \begin{cases} p-1, & \text{if } I_2 = \emptyset, \\ \lceil \frac{1}{2} \sum_{i \in I_2} \mathrm{con}(W_i) \rceil + |I_1|, & \text{otherwise.} \end{cases}$$

It is not hard to see that the partition inequalities (2.1) are valid for the ECON and NCON problems. The separation problem for partition inequalities is known to be NP-hard; see [GMS92]. However, there are fast heuristics for the separation of partition inequalities and our computational experiments have revealed that partition inequalities are very helpful for solving network connectivity problems. We know of no general necessary and sufficient conditions for partition inequalities to define facets. Some special cases are dealt with in

[GM90] and [S91].

We next consider the class of node partition inequalities which generalize the node cut inequalities (1.1)(ii) in a manner similar to how the partition inequalities (2.1) generalize the cut inequalities (1.1)(i). Consider a graph $G = (V, E)$ and requirement vector r. Let $Z_2 \subseteq \ldots \subseteq Z_k \subseteq V$, $k \geq 2$ be node sets with $|Z_j| = j - 1$ for $j = 2, \ldots, k$. Let $\{W_1, \ldots W_p\}$ be a partition of $V \backslash Z_k$, with $r(W_i) \geq 1$ for $i = 1, \ldots, p$, such that at least two node sets in the partition contains nodes of largest type k. Define $I_j := \{\, i \mid r(W_i) \geq j \,\}$ for $j = 1, \ldots, k$. The **node partition inequality** induced by $W_1, \ldots, W_p$ and $Z_2, \ldots, Z_k$ is given by

$$
\begin{aligned}
(2.2) \qquad x([W_1 : \ldots : W_p : Z_k]) \;&-\; x(\textstyle\bigcup_{j=2}^{k} \bigcup_{i \in I_j} [Z_j : W_i]) \\
&\geq\; p - 1.
\end{aligned}
$$

It is not difficult to show that the node partition inequalities are valid for the NCON problem. It is known to be NP-hard to separate this class of inequalities; see [GM90]. Some necessary and some sufficient conditions for these inequalities to define facets are known only in a few special cases; see [GMS89] and [S91].

A nice combinatorial relaxation of the ECON problem is the r-cover problem that can be defined as follows. Given a graph $G = (V, E)$ and positive integers r_v for all $v \in V$, an r**-cover** is a set $F \subseteq E$ of edges such that $|F \cap \delta(v)| \geq r_v$ for all $r \in V$. Cleary, every solution of the ECON problem defined by a graph G and node types $r \in \mathbf{N}^V$ is an r-cover. Edmonds' [E65] blossom inequalities for the 1-capacitated b-matching polytope of a graph $G = (V, E)$ can be transformed to the r-cover case to yield the r-**cover inequalities**, valid for the ECON problem, that have the following form:

$$x(E(H)) + x(\delta(H)\backslash T) \geq \lfloor (\textstyle\sum_{v \in H} r_v - |T|)/2 \rfloor \quad \text{for all } H \subseteq V, \text{ and all } T \subseteq \delta(H). \tag{2.3}$$

The separation problem for the class of r-cover inequalities can be solved in polynomial time using the separation algorithm of [PR82] for Edmond's blossom inequalities. In case some of the nodes have type 1, the r-cover inequalities can be strengthened as follows:

$$x(E(H))+x(\delta(H)\backslash T) \geq \lceil (\sum_{\substack{v \in H, \\ r_v \geq 2}} r_v - |T|)/2 \rceil + |\{ v \in H \mid r_v = 1 \}|. \tag{2.4}$$

These inequalities are valid for the ECON problem but not for the r-cover polytope. We solve the separation problem for the class of strengthened r-cover inequalities (2.4) heuristically.

3 Computational Results

At present, we have a preliminary version of a code for solving survivability problems with higher connectivity requirements. We first report about our computational results on random problems. We used the same set of random data as Ko and Monma [KM89] used for their high-connectivity heuristics so we will be able to compare results later. The test set consists of five complete graphs of 40 nodes and five complete graphs of 20 nodes, whose edge costs are independently drawn from a uniform distribution of real numbers between 0 and 20. For each of these 10 graphs, a minimum-cost k-edge connected subgraph for $k = 3, 4, 5$ is to be found. The table below reports the number of iterations (minimum and maximum) and the average time (in seconds on a SUN 4/50 IPX, a 28.5-MIPS machine) taken by our code to solve these problems for $k = 3, 4$, and 5, respectively. Only the time for the cutting plane phase is given.

# Nodes	**# Iterations**			**Average Time (secs)**		
$K =$	3	4	5	3	4	5
20 nodes	1-2	1-5	1-4	0.43	0.51	0.58
40 nodes	1-2	1-2	1-4	1.54	1.95	2.36

These excellent results were surprising, because we always thought high-connectivity problems to be harder than low-connectivity problems. But this does not seem to be true for random costs. The high-connectivity heuristics of Ko and Monma performed reasonably well. The relative gap between the heuristic (h) and the optimal solution value (o), namely $100 \times (h - o)/o$, computed for the above set of random problems, ranged between 0.8 and 12.8 with an average of 6.5 % error (taken over all problems).

One real-world application of survivable network design, where connectivities higher than two are needed, is the design of a fiber communication network that connects locations on a military ship containing various communication systems. The reason for demanding higher survivability of this network is obvious. The problem of finding a high-connected network topology minimizing the cable installation cost can be formulated as an NCON problem. We will describe the characteristics of this problem in the following. We obtained the graph and edge cost data of a generic ship model. It has the following features. The graph of possible link installations has the form of a three-dimensional grid with 15 layers, 494 nodes, and 1096 edges. Of the grid's 494 nodes, there are 461 regular nodes, 30 special nodes in the main part of the ship, and 3 priority nodes in the ship's tower. The notation "shipxyz" will be used to indicate that the regular nodes are of type x, special nodes are of type y, and priority nodes are of type z. So "ship013" is the problem, where the three nodes in the tower are of type 3, the 30 special nodes in the body of the ship are of type 1, and all other 461 grid nodes are of type 0. The cost structure is highly regular. So the problem is highly degenerate. Degeneracy together with the size of the ship problem caused us some difficulties. We were only able to obtain optimal solutions in two cases after considerable computational efforts.

Table 3.1 gives some computational results of our cutting plane algorithm on the several versions of the ship problem. The entries from left to right are:

PROBLEM	Problem name
IT	Number of iteration (i.e., LPs solved)
PART	Number of partition inequalities (2.1) added
RCOV	Number of r-cover inequalities (2.4) added
LB	Lower bound (i.e., optimal LP value)
GAP	(UB-LB)/LB in percent
TIME	in minutes:seconds

PROBLEM	VAR	IT	PART	RCOV	LB	UB	GAP	TIME
ship013	1088	3252	777261	0	211957.1	217428	2.58	10122:35
ship023	1088	15	4090	0	286274.0	286274	0	27:20
ship033	1082	42	10718	1	461590.6	483052	4.64	55:26
ship113	1090	128	17199	0	902691.0	918691	1.77	4724:55
ship123	1088	61	13210	0	906691.0	930691	2.57	1167:37
ship133	1084	176	21564	0	945052.0	1008808	6.74	119:15
ship223	1085	5	541	0	940925.0	940925	0	0:43
ship233	1081	5	532	0	1028193.0	1029176	0.09	0:54

Table 3.1: Performance of cutting plane algorithm on ship problems

Summarizing our computational results, we can say that for survivability problems with high connectivity requirements, many nodes of type 0 and highly regular cost structure (such as the ship problems) much still remains to be done to speed up our code and enhance the quality of solutions. This is in contrast to our previous work (see [GMS92]) on applications in the area of telephone network design, where we were able to obtain optimal solutions in a few minutes. However, these preliminary results provide hope for similar success in the future for these difficult network design problems. Further research on the polyhedral structure of these problems is needed along with algorithmic advances. Some steps in this direction have been taken by [C90] and [C91].

References

[CMW89] R.H. Cardwell, C.L. Monma, and T-H. Wu. Computer-aided design procedures for survivable fiber optic networks. *IEEE Journal on Selected Areas in Communications 7*, 1989.