

Chapter 10

Design of Survivable Networks

M. Grötschel

Konrad-Zuse-Zentrum für Informationstechnik Berlin, Heilbronner Str. 10, D-10711 Berlin, Germany

C.L. Monma

Bell Communications Research, 445 South Street, Morristown, NJ 07960, U.S.A

M. Stoer

Telenor Research, P.O. Box 83, N-2007 Kjeller, Norway

1. Overview

This chapter focuses on the important practical and theoretical problem of designing survivable communication networks, i.e., communication networks that are still functional after the failure of certain network components. We motivate this topic in Section 2 by using the example of fiber optic communication network design for telephone companies. A very general model (for undirected networks) is presented in Section 3 which includes practical, as well as theoretical, problems, including the well-studied minimum spanning tree, Steiner tree, and minimum cost k -connected network design problems.

The development of this area starts with outlining structural properties in Section 4 which are useful for the design and analysis of algorithms for designing survivable networks. These lead to worst-case upper and lower bounds. Heuristics that work well in practice are also described in Section 4. Polynomially-solvable special cases of the general survivable network design problem are summarized in Section 5.

Section 6 contains polyhedral results from the study of these problems as integer programming models. We present large classes of valid and often facet-defining inequalities. We also summarize the complexity of the separation problem for these inequalities. Finally we provide complete and nonredundant descriptions of a number of polytopes related to network survivability problems of small dimensions.

Section 7 contains computational results using cutting plane approaches based on the polyhedral results of Section 6 and the heuristics described in Section 4. The results show that these methods are efficient and effective in producing optimal or near-optimal solutions to real-world problems.

A brief review of the work on survivability models of directed networks is given in Section 8. We also show here how directed models can help to solve undirected cases.

2. Motivation

In this section we set the stage for the topic of this chapter by considering an application to designing communication networks for telephone companies based on fiber optic technology. We will use this to introduce the concept of survivability in network design and to motivate the general optimization models described in the next section. It will become clear later that our models capture many other situations that arise in practice and in theory as well.

Fiber optic technology is rapidly being deployed in communication networks throughout the world because of its nearly-unlimited capacity, its reliability and cost-effectiveness. The high capacity of new technology fiber transmission systems has resulted in the capability to carry many thousands of telephone conversations and high-speed data on a few strands of fiber. These advantages offer the prospect of ushering in many new information networking services which were previously either technically impossible or economically infeasible.

The economics of fiber systems differ significantly from the economics of traditional copper-based technologies. Copper-based technologies tend to be bandwidth limited. This results in a mesh-like network topology, which necessarily has many diverse paths between any two locations with each link carrying only a very small amount of traffic. In contrast, the high-capacity fiber technologies tend to suggest the design of sparse 'tree-like' network topologies. These topologies have only a few diverse paths between locations (often just a single path) and each link has a very high traffic volume. This raises the possibility of significant service disruptions due to the failure of a single link or single node in the network. The special report 'Keeping the phone lines open' by Zorpette [1989] describes the many man-made and natural causes that can disrupt communication networks, including fires, tornados, floods, earthquakes, construction or terrorist activities. Such failures occur surprisingly frequently and with devastating results as described in this report and in the popular press [e.g., see Newark Star Ledger, 1987, 1988a, 1988b; New York Times, 1988, 1989; Wall Street Journal, 1988].

Hence, it is vital to take into account such failure scenarios and their potential negative consequence when designing fiber communication networks. Recall that one of the major functions of a communication network is to provide connectivity between users in order to provide a desired service. We use the term 'survivability' to mean the ability to restore network service in the event of a catastrophic failure, such as the complete loss of a transmission link or a facility switching node. Service could be restored by means of routing traffic around the damage through other existing facilities and switches, if this contingency is provided for in the network architecture. This requires additional connectivity in the network topology and a means to automatically reroute traffic after the detection of a failure.

A network topology could provide protection against a single link failure if it remains connected after the failure of any single link. Such a network is called 'two-edge connected' since at least two edges have to be removed in order to disconnect the network. However, if there is a node in the network whose removal

does disconnect the network, such a network would not protect against a single node failure. Protection against a single node failure can be provided in an analogous manner by 'two-node connected' networks.

In the case of fiber communication networks for telephone companies, two-connected topologies provide an adequate level of survivability since most failures usually can be repaired relatively quickly and, as statistical studies have revealed, it is unlikely that a second failure will occur in their duration. However, for other applications it may be necessary to provide higher levels of connectivity.

One simple and cost effective means of using a two-connected topology to achieve an automatic recovery from a single failure is called diverse protection routing. Most fiber transmission systems employ a protection system to back up the working fiber systems. An automatic protection switch detects failure of a working system and switches the working service automatically to the protection system. By routing the protection system on a physically diverse route from the working system, one provides automatic recovery from a single failure at the small cost of a slightly longer path for the protection system than if it were routed on the same physical path as the working system. One would suspect, and studies have proven, that the additional cost of materials and installation of diverse fiber routes would be acceptable for the benefits gained [see Wu, Kolar & Cardwell, 1988; Wu & Cardwell, 1988; Kolar & Wu, 1988; and Cardwell, Wu & Woodall, 1988].

We consider the problem of designing minimum cost networks satisfying certain connectivity requirements that model network survivability. A formal model will be described in the next section. This model allows for different levels of survivability that arise in practical and theoretical models including the minimum spanning tree, Steiner tree, and minimum cost k -connected network design problems.

The application described here requires three distinct types of nodes: special offices, which must be protected from single edge or node failures, ordinary offices, which need only be connected by a single path, and optional offices, which may be included or excluded from the network design depending only upon cost considerations. The designation of office type is performed by a planner based on a cost/benefit analysis. Normally, the special offices are highly important and/or high-revenue-producing offices, with perhaps a high proportion of priority services. It may not be economically possible to ensure the service of ordinary or optional offices in the face of potential failures. In fact, some offices only have one outlet to the network and so it would be technologically impossible to provide recovery if this path were blocked by a network failure.

We note that two-connected network topologies are cost effective. For example, a typical real-world problem which we solved has a survivable network with cost of only 6% above the minimum spanning tree; however, a single failure in the tree network could result in the loss of 33% of the total traffic while the survivable network could lose at most 3% from any single failure. We also note that the heuristic methods described in Section 4 and the polyhedral cutting plane methods described in Section 6 are efficient and very effective in generating optimal and near-optimal network topologies as we describe in Section 7.

We conclude this section by pointing out that the topology design problem considered here is just the first step in the overall design of fiber communications networks. There are other issues that need to be addressed once the topology is in place. For instance, demands for services are usually in units of circuits, called DS0 rate, or bundles of 24 circuits, called DS1 rate. Fiber optic transmission rates come in units of 28 DS1s or 672 DS0s, called DS3 rate. Hence, it is necessary to place multiplexing equipment to convert between the three rates, and to route these demands through the network. Another issue is that the fiber signals need to be amplified using repeater equipment if the signals travel beyond a given threshold distance.

Furthermore, the network is generally organized into a facility hierarchy. That is, offices are grouped together into clusters, with each cluster having one hub office to handle traffic between clusters. This grouping considers such factors as community-of-interest and geographic area. Groups of clusters can be grouped into sectors, with each sector having one gateway office, which is a hub building designated to handle inter-sector traffic. This hierarchy allows traffic to be concentrated into high capacity routes to central locations where the demands are sorted according to destination. This concept of hub routing has proven to give near-optimal results, see Wu & Cardwell [1988]. A two-connected network allows for dual homing, i.e., for the possibility of splitting the demand at an office between a home hub and foreign hub to protect against a hub failure.

Given the complexity of the overall design problem, these issues are generally handled in a sequential fashion with the topology question being decided first. This process is often iterated until an acceptable network design is obtained. In this chapter, we only deal with the network topology design aspect and not the multiplexing or bundling aspects. For an overview of this combined process and a description of a computer-based planning tool [Bellcore, 1988] incorporating all of these features, see Cardwell, Monma & Wu [1989].

3. Integer programming models of survivability

In this section, we formalize the undirected survivable network design problems that are being considered in this chapter. Variants that are based on directed networks will be briefly treated in Section 8.

3.1. The general model for undirected networks

A set V of *nodes* is given representing the locations of the offices that must be interconnected into a network in order to provide the desired services. A collection E of *edges* is also specified that represent the possible pairs of nodes between which a direct transmission 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. For our range of applications, loops are irrelevant. Thus we assume in the following that

graphs have no loops. Parallel transmission links occur in practice; therefore, our graphs may have parallel edges. For technical reasons to be explained later, we will however restrict ourselves in this paper to simple graphs (i.e., loopless graphs without parallel edges) when we consider node survivability problems and node connectivity.

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 survivability conditions, which we describe below, are satisfied. We note that the cost here represents setting up the topology for the communication network and includes placing conduits in which to lay the fiber cable, placing the cables into service, and other related costs. We do not consider costs that depend on how the network is implemented such as routing, multiplexing, and repeater costs. Although these costs are also important, it is (as mentioned in Section 2) usually the case that a topology is designed first and then these other costs are considered in a second stage of optimization.

If $G = (V, E)$ is a graph, $W \subseteq V$ and $F \subseteq E$, then we denote by $G - W$ and $G - F$ the graph that is obtained from G by deleting the node set W and the edge set F , respectively. For notational convenience we write $G - v$ and $G - e$ instead of $G - \{v\}$ and $G - \{e\}$, respectively. The difference of two sets M and N is denoted by $M \setminus N$.

For any pair of distinct nodes $s, t \in V$, an $[s, t]$ -path P is a sequence of nodes and edges $(v_0, e_1, v_1, e_2, \dots, v_{l-1}, e_l, v_l)$, where each edge e_i is incident with the nodes v_{i-1} and v_i ($i = 1, \dots, l$), where $v_0 = s$ and $v_l = t$, and where no node or edge appears more than once in P . A collection P_1, P_2, \dots, P_k of $[s, t]$ -paths is called *edge-disjoint* if no edge appears in more than one path, and is called *node-disjoint* if no node (except for s and t) appears in more than one path. In standard graph theory two parallel edges are not considered as node-disjoint paths. For our applications it is sensible to do so. However, this modification would lead to considering nonstandard variations of node connectivity, reformulations of Menger's theorem etc. In order not to trouble the reader with these technicalities we have decided to restrict ourselves to simple graphs when node-disjoint paths are treated. The results presented in this paper carry, appropriately stated, over to the case where parallel edges are considered as node-disjoint paths. This theory is developed in Stoer [1992].

Two different nodes s, t of a graph are called *k-edge* (resp. *k-node*) *connected* if there are k edge-disjoint (resp. node-disjoint) paths between s and t . A graph with at least two nodes is called *k-edge* or *k-node connected* if all pairs of distinct nodes of G are *k-edge* or *k-node connected*, respectively. For our purposes, the graph K_1 consisting of just one node is *k-edge* and *k-node connected* for every natural number k . For a graph $G \neq K_1$, the largest integer k such that G is *k-edge connected* (resp. *k-node connected*) is denoted by $\lambda(G)$ (resp. $\kappa(G)$) and is called the *edge connectivity* (resp. *node connectivity*) of G . An *articulation set* of a connected graph G is a set of nodes whose removal disconnects G , and an *articulation node* is a single node disconnecting G .

The survivability conditions require that the network satisfy certain edge and node connectivity requirements. To specify these, three nonnegative integers r_{st} , k_{st} and d_{st} are given for each pair of distinct nodes $s, t \in V$. The numbers r_{st} represent the *edge survivability requirements*, and the numbers k_{st} and d_{st} represent the *node survivability requirements*; this means that the network $N = (V, F)$ to be designed has to have the property that, for each pair $s, t \in V$ of distinct nodes, N must contain at least r_{st} edge-disjoint $[s, t]$ -paths, and that the removal of at most k_{st} nodes (different from s and t) from N must leave at least d_{st} edge-disjoint $[s, t]$ -paths. (Clearly, we may assume that $k_{st} \leq |V| - 2$ for all $s, t \in V$, and we will do this throughout this chapter). These conditions ensure that some communication path between s and t will survive a prespecified level of combined failures of both nodes and links. The levels of survivability specified depend on the relative importance placed on maintaining connectivity between different pairs of offices.

Given $G = (V, E)$ and $r, k, d \in \mathbf{Z}_+^{E \times V}$, extend the functions r and d to functions operating on sets by setting

$$\text{con}(W) := \max\{r_{st} \mid s \in W, t \in V \setminus W\} \quad (1)$$

and

$$d(Z, W) := \max\{d_{st} \mid s \in W \setminus Z, t \in V \setminus (Z \cup W)\} \text{ for } Z, W \subseteq V. \quad (2)$$

We call a pair (Z, W) with $Z, W \subseteq V$ *eligible* (with respect to k) if $Z \cap W = \emptyset$ and $|Z| = k_{st}$ for at least one pair of nodes s, t with $s \in W$ and $t \in V \setminus (Z \cup W)$.

Let us now introduce a variable x_e for each edge $e \in E$, and consider the vector space \mathbf{R}^E . Every subset $F \subseteq E$ induces an *incidence vector* $\chi^F = (\chi_e^F)_{e \in E} \in \mathbf{R}^E$ by setting $\chi_e^F := 1$ if $e \in F$, $\chi_e^F := 0$ otherwise; and vice versa, each 0/1-vector $x \in \mathbf{R}^E$ induces a subset $F^x := \{e \in E \mid x_e = 1\}$ of the edge set E of G . If we speak of the incidence vector of a path in the sequel we mean the incidence vector of the edges of the path. We can now formulate the network design problem introduced above as an integer linear program with the following constraints.

$$\begin{aligned} \text{(i)} \quad & \sum_{i \in W} \sum_{j \in V \setminus W} x_{ij} \geq \text{con}(W) \quad \text{for all } W \subseteq V, \emptyset \neq W \neq V, \\ \text{(ii)} \quad & \sum_{i \in W} \sum_{j \in V \setminus (Z \cup W)} x_{ij} \geq d(Z, W) \quad \text{for all eligible } (Z, W) \text{ of subsets of } V, \\ \text{(iii)} \quad & 0 \leq x_{ij} \leq 1 \quad \text{for all } ij \in E, \\ \text{(iv)} \quad & x_{ij} \text{ integral} \quad \text{for all } ij \in E. \end{aligned} \quad (3)$$

Note that if $N - Z$ contains at least d_{st} edge-disjoint $[s, t]$ -paths for each pair s, t of distinct nodes in V and for each set $Z \subseteq V \setminus \{s, t\}$ with $|Z| = k_{st}$, and if $r_{st} = k_{st} + d_{st}$, then all node survivability requirements are satisfied, i.e., inequalities of type (3ii) need not be considered for node sets $Z \subseteq V \setminus \{s, t\}$ with $|Z| < k_{st}$. It follows from Menger's theorem (see [Frank, 1995]) that, for every feasible solution x of (3), the subgraph $N = (V, F^x)$ of G defines a network that satisfies the given edge and node survivability requirements.

This model, introduced in Grötschel & Monma [1990], generalizes and unifies a number of problems that have been investigated in the literature either from a practical or theoretical point of view. We mention here some of these cases.

The classical *network synthesis problem* for multiterminal flows (see Chapter 13) is obtained from (3) by dropping the constraints (3ii) and (3iv). In the standard formulation of the network synthesis problem, the upper bounds $x_e \leq 1$ are not present. But our model allows parallel edges in the underlying direct-link graph, or equivalently, allows the upper bound in constraints (3iii) to take on any nonnegative values for each edge. This linear programming problem has a number of constraints that is exponential in the number of nodes of G . For the case $c_{ij} = c$ for all $ij \in E$, where c is a constant, Gomory & Hu [1961] found a simple algorithm for its solution. Bland, Goldfarb & Todd [1981] pointed out that the separation problem for the class of inequalities (3i) can be solved in polynomial time by computing a minimum capacity cut; thus, it follows by the ellipsoid method that the classical network synthesis problem can be solved in polynomial time. (See Grötschel, Lovász & Schrijver [1988] for details on the ellipsoid algorithm and its applications to combinatorial optimization.) The *integer network synthesis problem*, i.e., the problem obtained from (3) by dropping constraints (3ii) and the upper bounds $x_e \leq 1$, was solved by Chou & Frank [1970], for the case $c_{ij} = c$ for all $ij \in V \times V$, and $r_{ij} \geq 2$ for all ij .

The *minimum spanning tree problem* can be phrased as the task to find a minimum-cost connected subset $F \subseteq E$ of edges spanning V (see Chapter 12). This problem can be viewed as a special case of (3) as follows. We drop the constraints (3ii) and set $r_{st} := 1$ for all distinct $s, t \in V$ in constraints (3i).

Similarly, the closely related *Steiner tree problem* is to find a minimum-cost connected subset $F \subseteq E$ of edges that span a specified subset $S \subseteq V$ of nodes. This problem is a special case of (3) where we drop constraints (3ii) and set $r_{st} := 1$ in constraints (3i), for all $s, t \in S$, and $r_{st} := 0$ otherwise. (See also Chapter 12.) Let us remark at this point that the Steiner tree problem is well-known to belong to the class of NP-hard problems. As it is a special case of (3), our general problem of designing survivable networks is NP-hard as well. (See Garey & Johnson [1979] for details on the theory of NP-completeness.)

The problem of finding a *minimum-cost k -edge connected network* in a given graph is a special case of (3) where all inequalities (3ii) are dropped and where $r_{st} = k$ for all distinct $s, t \in V$. The problem of finding an *optimal k -node connected network*, for $k \leq |V| - 1$, is a special case of (3) where we drop the constraints (3i) and set $k_{st} := k - 1$ and $d_{st} := 1$ for all distinct $s, t \in V$.

3.2. A brief discussion of the model

The rest of this chapter is mainly devoted to studying various aspects of model (3) and some of its special cases. Among other subjects, we describe heuristics to compute upper and lower bounds for the optimum value of (3). To compute a lower bound, one is naturally led to dropping the integrality constraints (3iv)

and solving the LP-relaxation (3i), (3ii) and (3iii) of the survivable network design problem.

Two questions immediately arise. Can one solve this LP despite the fact that it has exponentially many constraints? Can one find a better (or a series of increasingly better) LP-relaxations that are solvable in polynomial (or practically reasonable) time?

We will address the first question in Section 7 and the second in Section 6. But here we would like to give a glimpse at the approach that leads to answering these questions. The method involved is known as polyhedral combinatorics. It is a vehicle to provide (in some sense) the best possible LP-relaxation. We want to demonstrate now how the second question leads to the investigation of certain integral polyhedra in a very natural way. See Grötschel & Padberg [1985] and Pulleyblank [1989] for a survey on polyhedral combinatorics.

To obtain a better LP-relaxation of (3) than the one arising from dropping the integrality constraints (3iv), we define the following polytope. Let $G = (V, E)$ be a graph, let $E_V := \{st \mid s, t \in V, s \neq t\}$, and let $r, k, d \in \mathbb{Z}_+^{E_V}$ be given. Then

$$\text{CON}(G; r, k, d) := \text{conv}\{x \in \mathbb{R}^E \mid x \text{ satisfies (3i)–(3iv)}\} \quad (4)$$

is the polytope associated with the network design problem given by the graph G and the edge and node survivability requirements r, k , and d . (Above 'conv' denotes the convex hull operator.) In the sequel, we will study $\text{CON}(G; r, k, d)$ for various special choices of r, k and d . Let us mention here a few general properties of $\text{CON}(G; r, k, d)$ that are easy to derive.

Let $G = (V, E)$ be a graph and $r, k, d \in \mathbb{Z}_+^{E_V}$ be given as above. We say that $e \in E$ is *essential with respect to* $(G; r, k, d)$ (short: $(G; r, k, d)$ -essential) if $\text{CON}(G - e; r, k, d) = \emptyset$. In other words, e is essential with respect to $(G; r, k, d)$ if its deletion from G results in a graph such that at least one of the survivability requirements cannot be satisfied. We denote the set of edges in E that are essential with respect to $(G; r, k, d)$ by $\text{ES}(G; r, k, d)$. Clearly, for all subsets $F \subseteq E \setminus \text{ES}(G; r, k, d)$, $\text{ES}(G; r, k, d) \subseteq \text{ES}(G - F; r, k, d)$ holds. Let $\dim(S)$ denote the *dimension* of a set $S \subseteq \mathbb{R}^n$, i.e., the maximum number of affinely independent elements in S minus 1. Then one can easily prove the following two results [see Grötschel & Monma, 1990].

Theorem 1. *Let $G = (V, E)$ be a graph and $r, k, d \in \mathbb{Z}_+^{E_V}$ such that $\text{CON}(G; r, k, d) \neq \emptyset$. Then*

$$\text{CON}(G; r, k, d) \subseteq \{x \in \mathbb{R}^E \mid x_e = 1 \text{ for all } e \in \text{ES}(G; r, k, d)\}, \text{ and} \\ \dim(\text{CON}(G; r, k, d)) = |E| - |\text{ES}(G; r, k, d)|.$$

An inequality $a^T x \leq \alpha$ is *valid* with respect to a polyhedron P , if $P \subseteq \{x \mid a^T x \leq \alpha\}$; the set $F_a := \{x \in P \mid a^T x = \alpha\}$ is called the *face* of P defined by $a^T x \leq \alpha$. If $\dim(F_a) = \dim(P) - 1$ and $F_a \neq \emptyset$ then F_a is a *facet* of P , and $a^T x \leq \alpha$ is called *facet-defining* or *facet-inducing*.

Theorem 2. Let $G = (V, E)$ be a graph and $r, k, d \in \mathbf{Z}_+^{E_v}$ such that $\text{CON}(G; r, k, d) \neq \emptyset$. Then

- (a) $x_e \leq 1$ defines a facet of $\text{CON}(G; r, k, d)$ if and only if $e \in E \setminus \text{ES}(G; r, k, d)$;
- (b) $x_e \geq 0$ defines a facet of $\text{CON}(G; r, k, d)$ if and only if $e \in E \setminus \text{ES}(G; r, k, d)$ and $\text{ES}(G; r, k, d) = \text{ES}(G - e; r, k, d)$.

Theorems 1 and 2 solve the dimension problem and characterize the trivial facets. But these characterizations are (in a certain sense) algorithmically intractable as the next observation shows, which follows from results of Ling & Kameda [1987].

Remark 1. The following three problems are NP-hard.

Instance: A graph $G = (V, E)$ and vectors $r, k, d \in \mathbf{Z}_+^E$.

Question 1: Is $\text{CON}(G; r, k, d)$ nonempty?

Question 2: Is $e \in E$ ($G; r, k, d$)-essential?

Question 3: What is the dimension of $\text{CON}(G; r, k, d)$?

However, for most cases of practical interest in the design of survivable networks, the sets $\text{ES}(G; r, k, d)$ of essential edges can be determined easily, and thus the trivial LP-relaxation of (3) can be set up without difficulties by removing the redundant inequalities identified by Theorem 2.

3.3. A model used in practice

The model discussed so far mixes node and edge connectivity, and provides, for each pair of nodes, the possibility to specify particular connectivity requirements. It is, as mentioned before, a quite general framework that models many practical situations simultaneously. This generality demands a considerable amount of data. In our real-world application, it turned out that the network designers were either interested in node connectivity or in edge connectivity requirements but not in both simultaneously. Also, the data for implementing the general model were not available in practice. A specialized version, to be described below, proved to be acceptable from the point of view of data acquisition and was still considered a reasonable model of reality by practitioners.

To model these (slightly more restrictive) survivability conditions, 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. For any $W \subseteq V$, the integer $r(W) := \max\{r_v \mid v \in W\}$ is called the *type* of W . We say that the network $N = (V, F)$ to be designed satisfies the *node survivability conditions* if, for each pair $s, t \in V$ of distinct nodes, N contains at least

$$r_{st} := \min\{r_s, r_t\} \tag{5}$$

node-disjoint $[s, t]$ -paths. Similarly, we say that $N = (V, F)$ satisfies the *edge survivability conditions* if, for each pair $s, t \in V$ of distinct nodes, N contains r_{st} edge-disjoint $[s, t]$ -paths. These conditions ensure that some communication

path between s and t will survive a prespecified level of node or link failures. We will discuss these special cases in more detail later. To make the (somewhat clumsy) general notation easier, we introduce further symbols and conventions to denote these node- or edge-survivability models. Given a graph $G = (V, E)$ and a vector of node types $r = (r_s)_{s \in V}$ we assume — without loss of generality — that there are at least two nodes of the largest type. If we say that we consider the k NCON problem (for G and r) then we mean that we are looking for a minimum-cost network that satisfies the node survivability conditions and where $k = \max\{r_s \mid s \in V\}$. Similarly, we speak of the k ECON problem (for G and r), when we consider the edge survivability conditions. When we want to leave it open whether the survivability problem is specified by node types or by a $V \times V$ matrix r (k and d), we speak of ECON (resp., NCON) problems.

Let $G = (V, E)$ be a graph. For $Z \subseteq V$, let $\delta_G(Z)$ denote the set of edges with one end node in Z and the other in $V \setminus Z$. It is customary to call $\delta_G(Z)$ a *cut*. If it is clear with respect to which graph a cut $\delta_G(Z)$ is considered, we simply drop the index and write $\delta(Z)$. We also write $\delta(v)$ for $\delta(\{v\})$. If X, Y are subsets of V with $X \cap Y = \emptyset$, we set $[X : Y] := \{ij \in E \mid i \in X, j \in Y\}$, thus $\delta(X) = [X : V \setminus X]$. 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 $s \in V$ [using (5) in the definition of $\text{con}(W)$ in (1)]:

$$\begin{aligned} & \min \quad c^T x \\ & \text{subject to} \\ & \quad \text{(i) } x(\delta(W)) \geq \text{con}(W) \text{ for all } W \subseteq V, \emptyset \neq W \neq V; \\ & \quad \text{(ii) } x(\delta_{G-Z}(W)) \geq \text{con}(W) - |Z| \\ & \hspace{15em} \text{for all pairs } s, t \in V, s \neq t, \text{ and} \quad (6) \\ & \hspace{15em} \text{for all } \emptyset \neq Z \subseteq V \setminus \{s, t\} \text{ with } |Z| \leq r_{st} - 1, \\ & \hspace{15em} \text{and for all } W \subseteq V \setminus Z \text{ with } s \in W, t \notin W; \\ & \quad \text{(iii) } 0 \leq x_{ij} \leq 1 \quad \text{for all } ij \in E; \\ & \quad \text{(iv) } x_{ij} \text{ integral} \quad \text{for all } ij \in E. \end{aligned}$$

It follows from Menger's theorem that the feasible solutions of (6) are the incidence vectors of edge sets F such that $N = (V, F)$ satisfies all node survivability conditions; i.e., (6) is an integer programming formulation of the k NCON problem. Deleting inequalities (6ii) we obtain, again from Menger's theorem, an integer programming formulation for the k ECON problem. The inequalities of type (6i) will be called *cut inequalities* and those of type (6ii) will be called *node cut inequalities*.

The polyhedral approach to the solution of the k NCON (and similarly the k ECON) problem consists of studying the polyhedron obtained by taking the convex hull of the feasible solutions of (6). We set

$$\begin{aligned} k\text{NCON}(G; r) &:= \text{conv}\{x \in \mathbf{R}^E \mid x \text{ satisfies (6i)–(6iv)}\}, \\ k\text{ECON}(G; r) &:= \text{conv}\{x \in \mathbf{R}^E \mid x \text{ satisfies (6i), (6iii) and (6iv)}\}. \end{aligned}$$

To tie this notation with the previously introduced more general concept, note that

$$k\text{ECON}(G; r) = \text{CON}(G; r', 0, 0),$$

where $r' \in \mathbf{R}^{V \times V}$ with $r'_{st} := \min\{r_s, r_t\}$ for all $s, t \in V$. Also, if there are no parallel edges then

$$k\text{NCON}(G; r) = \text{CON}(G; r', k', d'),$$

where $k'_{st} := \max\{0, r'_{st} - 1\}$ for all $s, t \in V$ and $d' := r' - k'$.

This survivability model, the polyhedra, and the integer and linear programming problems associated with the $k\text{ECON}$ and $k\text{NCON}$ problems, will be studied in more detail in the sequel.

4. Structural properties and heuristics

In this section we describe some heuristic approaches for the solution of ECON and NCON problems. There are standard methods like greedy and interchange heuristics and heuristics that are motivated by techniques for the approximate solution of other NP-hard problems, like the (much investigated) traveling salesman problem. Some heuristics are more special-purpose since they make use of structural properties of k -connected graphs. A few structural results and their uses are reviewed in Sections 4.1 and 4.2. Section 4.2 concentrates on the design of practically-effective heuristics for ECON and NCON problems, while Section 4.3 discusses lower bounds and heuristics with worst-case performance guarantees.

4.1. Lifting and the structure of optimum solutions

Connectivity is a very rich and active topic of graph theory and it is conceivable that the knowledge that has been accumulated in this area can be exploited further for the design of effective approximation algorithms. We do not attempt to cover the work on connectivity here in detail and refer to Frank [1995] for a comprehensive survey of connectivity results. We will just mention a few structural properties of connected graphs that have been employed for the design of practically-useful heuristics. We begin by describing a ‘local’ construction technique. Lovász [1976] & Mader [1978] [see also Frank, 1992b] have proved so-called ‘lifting theorems’ that show that simple manipulations of a graph can be made without destroying certain edge-connectivity properties. These manipulations are useful for construction heuristics for the ECON problem (where parallel edges are allowed) for general connectivity requirements r_{vw} .

In order to state these results we have to introduce a few definitions. Let $G = (V, E)$ be a graph, and let x be a node of G that we call *special*. We assume that x is adjacent to distinct nodes y and z . The graph G' obtained from G by deleting the edges xy and xz and adding the edge yz is called *lifting* of G at x .

If the edge-connectivity of G' between any two nodes of $V' - x$ is not smaller than that of G , then the lifting is called *admissible*.

Theorem 3 [Mader, 1978]. *Let $G = (V, E)$ be a graph with special node x .*

(a) *If the degree of x is at least 4 and x is not an articulation node then G has an admissible lifting at x .*

(b) *If x is an articulation node and no single edge incident to x is a cut then G has an admissible lifting at x .*

Theorem 4 [Lovász, 1976]. *Let $G = (V, E)$ be a Eulerian graph, let x be a (special) node of even degree, and let y be any node adjacent to x in G . Then there is another neighbor node z of x , such that the lifting at x involving y and z is admissible.*

Consider the k ECON and k NCON problems for the complete graph K_n where $r_x = k$ for all nodes x , and where the nonnegative costs satisfy the triangle inequality, i.e., $c_{xz} \leq c_{xy} + c_{yz}$ for all nodes x, y and z . Using the Lifting Theorems 3 and 4, Monma, Munson & Pulleyblank [1990] showed the following result for k ECON for K_n with $k = 2$.

Theorem 5. *Given costs satisfying the triangle inequality, there is an optimal k -edge connected spanning network satisfying the following conditions:*

- (a) *all nodes are of degree k or $k + 1$; and*
- (b) *removing any set of at most k edges does not leave all of the resulting connected components k -edge connected.*

It is easy to see that if the cost function satisfies the triangle inequalities, there is an optimal 2-node connected solution with cost equal to an optimal 2-edge connected solution. In fact, Monma, Munson & Pulleyblank [1990] show that the term ' k -edge' can be replaced by ' k -node' throughout Theorem 5 to obtain a similar result for k NCON for K_n with $k = 2$. Furthermore, they show that these 'characterize' the optimal 2-connected networks in the following sense: given any 2-connected graph $G = (V, E)$ satisfying conditions (a) and (b) of Theorem 5 for $k = 2$, then there exist costs satisfying the triangle inequality such that G is the unique optimal solution.

Bienstock, Brickell & Monma [1990] showed that Theorem 5 holds for arbitrary k . They also showed that ' k -edge' can be replaced by ' k -node' throughout Theorem 5 with the technical restriction that $|V| \geq 2k$ in order for (a) to hold. This technical restriction is necessary since without it there is an infinite family of examples where condition (a) fails.

We note that the cost of an optimal k -edge connected solution may be strictly less than the cost of an optimal k -node connected solution for any $k \geq 3$, and that the conditions in Theorem 5 do not characterize the optimal solutions as they do in the case $k = 2$. The proof of Theorem 5 for the k -edge connected case uses the Lifting Theorem 3. The proof of Theorem 5 in the k -node connected case is

much more difficult and requires the use of pairs of liftings as well as a number of further technical results.

4.2. Construction and improvement heuristics

We now describe some heuristics for constructing feasible networks and heuristics for improving the cost of a feasible solution for the k ECON and k NCON problems. This is, to a large extent, a summary of the work of Monma & Shallcross [1989] for the ‘low-connected’ survivable network design problem, that is, problems with node types in $\{0, 1, 2\}$. The performance on real-world problems is described in Section 7. These heuristics were inspired by the wide variety of heuristics for the traveling salesman problem and other combinatorial optimization problems and were designed to work on sparse underlying graphs. The heuristics are used in a local search approach to obtain low-cost network designs; see Papadimitriou & Steiglitz [1982] for a general discussion of local search procedures for combinatorial optimization problems. It is obvious how these heuristics can be generalized and applied to the general survivable network design problem, and so we just briefly mention one such extension.

One useful structural fact is that any edge-minimal two-connected graph $G = (V, E)$ can be constructed by an ‘ear decomposition’ procedure; see Lovász & Plummer [1986]. That is, first find a cycle C in G . Then repeatedly find a path P , called an ear, that starts at a node v in the solution, passes through nodes not yet in the solution, and ends up at a node w in the solution. All edge-minimal two-edge connected graphs can be constructed in this manner. If the nodes v and w are required to be distinct, then every edge-minimal two-node connected graph can be constructed in this manner.

The ear decomposition approach can be employed to construct a feasible 2-connected subgraph of a graph $G = (V, E)$ using costs to add ears in a greedy fashion. We call this the *greedy ears construction heuristic*. The first step is to construct a partial solution consisting of a cycle C spanning the set of nodes of type 2, called the ‘special nodes’. This is done by randomly selecting a special node v , and then selecting a special node w whose shortest path P to v is the longest among all special nodes. Let node u be the node next to w on the path P . We will construct a short cycle through the edge uw by finding a shortest path from u to w not using the edge uw . (There must be such a path; if not, there would not two node-disjoint paths between the special nodes v and w and so the problem would be infeasible). The next step is to repeatedly add ‘short’ ears to the current partial solution until all special nodes are on this two-connected network. This is done by first selecting a special node z , not yet in the solution, whose shortest path P to the partial solution is longest among all special nodes not yet included. We will find another shortest path Q from z to the partial solution that does not use any edges of P and that terminates on the partial solution at a node w other than v . (Again, such a path must exist for the problem to be feasible). The combination of paths P and Q must contain an ear, which is added to the partial solution.

A second construction heuristic uses the ear decomposition approach to construct a feasible 2-connected subgraph of a graph $G = (V, E)$ in a random fashion. We call this the *random sparse construction heuristic*. The first step is to construct an initial random cycle C spanning a subset of the special nodes. This is done by randomly choosing a special node v , and constructing a depth-first-search tree T rooted at v . Form a cycle by randomly choosing an edge of the form vz that is not in T . (There must be such an edge or else v is not on any cycle and so the problem is infeasible.) Next, random ears are repeatedly added until all special nodes are on the two-connected network. This is done by first constructing a depth-first-search forest F rooted at the nodes that are in the partial solution. A node v is said to be *allowed* if v is not yet in the solution, but it has an edge vw in E but not F , where w is in the solution. Randomly choose a node v from among the allowed nodes. (Again, there must be such a node or the problem is infeasible.) Let T be the tree in the forest F containing v , and let z be the root of T . The random ear chosen is the path from v to z in T , together with the edge vw . Since this method does not use cost information, generally it does not produce a low cost solution. However, this method is useful for generating starting random initial solutions on which to apply the improvement methods.

Improvement heuristics apply local transformations to any feasible network in order to reduce the cost while preserving feasibility. These transformations are applied until a locally optimal network is obtained; that is, no further improvements of these types are possible. Six local transformations are described in the sequel. These transformations are general enough to cover a wide range of feasible topologies, yet fast enough to be performed in real time, even on a personal computer.

Every two-connected network contains at least one cycle, and often is made up of many interconnected cycles. Furthermore, replacing the edges in a cycle C by edges forming a cycle C' on the same nodes preserves the feasibility of a solution. So it is natural to draw upon the extensive research on the traveling salesman problem [see Lawler, Lenstra, Rinnooy Kan & Shmoys, 1985, or Reinelt, 1994, or Chapter 3 of this handbook] for finding a near-optimal cycle. This is the basis of the two-optimal cycle and three-optimal cycle improvement heuristics. The pretzel, quetzel and degree improvement heuristics alter the structure of the two-connected part of the solution in less obvious ways. The one-optimal improvement heuristic alters the structure of the entire solution. These heuristics are described below.

The *two-optimal interchange heuristic* attempts to replace two edges of a cycle C by two edges not in the cycle to form a new cycle C' of lower cost. Similarly, the *three-optimal interchange heuristic* attempts to replace three edges of a cycle C by three edges not in the cycle to form a new cycle C' of lower cost.

These improvement heuristics replace one cycle C by another cycle C' on the same nodes, and so do not change the fundamental structure underlying the solution. The *pretzel transformation* replaces an edge uv of a cycle C by two crossing edges ux and vy to form a 'pretzel' P , where the nodes u, v, x and y appear in order on the cycle.

The *quetzel transformation* is the reverse of the pretzel transformation; that is, a

pretzel P is replaced by a cycle C by removing two crossing edges ux and vy and adding an edge uv .

As mentioned before, if the costs satisfy the triangle inequality, then there is an optimal two-connected solution where all nodes are of degree two or three. The proof of this result employs the Lifting Theorems 3 and 4. We will describe now the algorithmic use of these theorems in the form of *degree improvement transformations*.

Let node u be of degree four or more, and let nodes a, b, c , and d be four of its neighbors. There are three cases to consider. In Case a, there are node-disjoint paths P_{ab} and P_{bc} from a to b , and from b to c , respectively, which miss node u . In this case, edge ub is a chord and can be deleted. So we may assume that no three of the nodes a, b, c and d have such paths. Therefore, the paths P_{ab} and P_{bc} must intersect in node v and the paths P_{bc} and P_{cd} must intersect in node w . If nodes v and w are different, we are in Case b, and edges bu and uc are removed, and edge bc is added; this preserves the connectivity requirements and does not increase the cost if the triangle inequality holds. If nodes v and w are the same node, we are in Case c. Let a', b', c' and d' be the neighbors of $v(=w)$. Edges $au, ub, b'v$ and vc' are removed, and edges ab and $b'c'$ are added; this preserves the connectivity requirements and does not increase the cost if the triangle inequality holds. In all cases, the degree of node u is decreased, and no other degrees are increased; so repeated application of these transformations guarantees that an optimal solution where all degrees are two or three will be obtained so long as the triangle inequality holds.

All of the previous improvement heuristics operated only on the two-connected part of the solution. The *one-optimal interchange heuristic* considers the entire solution. This heuristic attempts to remove an edge uv from the current feasible solution and replace it with another edge of the form ux not currently in the solution. Such an interchange is made only if the resultant network is feasible and of lower cost.

These heuristics were tested on randomly generated problems and on real-world problems of telephone network design; see Monma & Shallcross [1989] for details. They could decrease the cost of manually constructed solutions by about 10% in a test of these methods on a real-world application. Comparison with the optimal solutions computed by the cutting plane algorithm (to be described in Section 7) on the same examples show that the gap was usually very small, about 1% of the optimal value. The highest running time for a sparse 116-node problem was 156 seconds on an IBM-PC/AT.

Ko & Monma [1989] modified the low-connectivity heuristics of Monma and Shallcross to the design of k -edge or k -node connected networks. A first feasible solution is constructed either by deleting edges successively from the whole graph while maintaining feasibility, or by adding successively k edge-disjoint $[1, j]$ -paths of minimum overall length, for all nodes $j \neq 1$. The output is necessarily k -edge connected. The local transformations for the low-connectivity case carry over to the high-connectivity case, except that the feasibility checks have to be implemented differently.

These heuristics could only be tested on random examples, as real-world examples were not yet available. When the cost of the best heuristic solution was compared with the optimal solution produced by our cutting plane algorithm, the gap was approximately 6%. Running times on a VAX 8650 ranged between 13 s for dense graphs of 20 nodes and $k = 3$ and 120 s for dense graphs of 40 nodes and $k = 5$.

Let us remark that to our knowledge the first heuristics for the design of minimum-cost survivable networks under general connectivity requirements date back to Steiglitz, Weiner & Kleitman [1969]. Their heuristic consists of a randomized starting routine and an optimization routine where local transformations are applied to a feasible solution as follows. Given a random ordering of the nodes, the *starting routine* adds edges between the first node with the highest connectivity requirement and the first node with the next highest connectivity requirement. In each step the connectivity requirements are updated. If the solution is feasible, the *optimizing routine* tries to improve this solution by successively replacing one pair of edges with another pair of edges to obtain another feasible solution of lower cost until no more improvements can be made this way.

They applied their heuristics to two real-world problems with 10 nodes and 58 nodes, and connectivity requirements in $\{3, 4, 5, 6\}$ and $\{6\}$, respectively. (Unfortunately, the data are not available any more.) The 58-node problem took about 12 minutes (on a UNIVAC 1108) per local optimum. Since no lower bounds on the optimal value for these problems are given, we cannot say how well these heuristics work.

4.3. Heuristics with performance guarantees

The last remark leads us to an important issue: quality guarantees, i.e., worst-case performance and lower bounds. The polyhedral approach, to be described later, can be viewed as a technique to obtain very good lower bounds for the optimum value of a k ECON or k NCON problem. But sometimes nice performance guarantees for heuristics or estimates for optimum values can be given with less elaborate techniques.

Let us first relate the 2-edge connectivity problems to the traveling salesman problem. Since every Hamiltonian cycle is 2-node (and thus 2-edge) connected, the optimum TSP-value, CTSP say, is not smaller than the optimum value COPT of the 2ECON problem with node types $r_v = 2$ for all v . On the other hand, using Theorem 1, Monma, Munson & Pulleyblank [1990] were able to show that if the triangle inequality holds, COPT can be bounded by CTSP from below by a constant factor, more precisely

$$\frac{3}{4} \text{CTSP} \leq \text{COPT} \leq \text{CTSP}.$$

To solve the 2ECON problem approximately, Frederickson & Jájá [1982] modified the Christofides heuristic for the traveling salesman problem and proved that, when the triangle inequality holds, the solution attains a cost CCHR with the same worst-case bound as in the traveling salesman problem, namely

$$\text{CCHR} \leq \frac{3}{2} \text{COPT.}$$

Another type of lower bound for the k -edge connected network design problem can be derived from the subtour elimination polytope, which is a natural linear programming relaxation of the traveling salesman polytope. Let CSUB denote the value of an optimal solution to the subtour elimination linear program; see Chapter 3. Goemans & Bertsimas [1993] showed that

$$\frac{k}{2} \text{CSUB} \leq \text{COPT.}$$

For $k = 2$, this was previously shown by Cunningham, see Monma, Munson & Pulleyblank [1990]. These results make use of the Lifting Theorem 4.

For edge connectivity problems with varying edge connectivity requirements the following heuristics are known. Goemans & Bertsimas [1993] developed a tree heuristic with worst-case guarantee for a version of k ECON problem with general node types $r \in \mathbf{Z}_+^V$, where edges may be used more than once. This means that a feasible solution to this problem is a vector $x \in \mathbf{Z}_+^E$ of nonnegative integers that satisfies all cut inequalities of $x(\delta(W)) \geq r(W)$, but not necessarily the upper bounds $x_e \leq 1$. A component $x_e \geq 2$ may be interpreted as 'edge e used x_e times'.

Let $\rho_1 < \rho_2 < \dots < \rho_p$ be the ordering of the distinct node types in r , let $\rho_0 := 0$, and let V_k be the set of nodes of type at least ρ_k , $k = 1, \dots, p$.

Tree heuristic

Compute, for all pairs of nodes u, v , the shortest path length c'_{uv} with respect to the given costs c .

Set $x_e := 0$ for all $e \in V \times V$.

For $k = 1$ to p do

– compute $T_k = (V_k, E_k)$ as the minimum spanning tree of the complete graph induced by V_k with respect to costs c'_e ;

– set $x_e := x_e + (\rho_k - \rho_{k-1})$ for all $e \in E_k$.

For each edge $e = (u, v)$ with $c'_e < c_e$ and $x_e > 0$, decrease x_e to 0 and increase by x_e the weights on the edges of a shortest $[u, v]$ -path.

Output x_e for $e \in E$.

Note that $x \in \mathbf{Z}_+^E$ satisfies all cut inequalities $x(\delta(W)) \geq r(W)$. Let y be the solution to the LP-relaxation of the ECON problem consisting of cut inequalities (2.4i) and nonnegativity constraints $x_e \geq 0$. Goemans & Bertsimas show that

$$\frac{x^T c}{y^T c} \leq \left(2 - \frac{2}{|V_1|}\right) \left(\sum_{k=1}^p \frac{\rho_k - \rho_{k-1}}{\rho_k}\right).$$

This bound is tight, if we consider for instance a 2ECON problem on a cycle of costs 1. Goemans & Bertsimas also describe a more refined tree heuristic with a better worst-case guarantee.

Agrawal, Klein & Ravi [1991] state a heuristic for an ECON problem, where edge-connectivity requirements are given by $r \in \mathbf{Z}_+^{V \times V}$ and the use of multiple

parallel edges is allowed in the solution. They prove that their algorithm outputs a solution whose cost is approximately within $2 \log R$ of the optimal, where R is the highest requirement value.

The worst-case guarantees for the heuristics of Goemans & Bertsimas [1993] and of Agrawal, Klein & Ravi [1991] were found by reduction to an ECON problem with costs satisfying the triangle inequality (see Step 1 and 4 of the tree heuristic). This reduction does not work, if the use of parallel edges is forbidden in the solution. For the case that the use of parallel edges is forbidden, the edge connectivity requirements are given by a node type vector r , and the cost function is arbitrary (but nonnegative), Goemans, Mihail, Vazirani & Williamson [1992] proposed a heuristic based on a primal-dual approach (using the cut inequalities). This heuristic has an approximation factor of

$$2 \sum_{i=1}^p \mathcal{H}(\rho_i - \rho_{i-1}),$$

where \mathcal{H} is the harmonic function $\mathcal{H}(k) = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{k}$, and where ρ_i ($i = 1, \dots, p$) is defined as above. In particular, for k -edge connectivity problems, the approximation factor is $2 \ln k$, and for $k = 2$ it is 3.

For k -edge connectivity problems one can do even better. Khuller & Vishkin [1994] gave a simple heuristic for finding a minimum-cost k -edge connected subgraph of a graph (where parallel edges do not appear in the solution). This heuristic has a worst-case guarantee of 2, even when the costs do not satisfy the triangle inequality. Worst-case guarantees for heuristics for the NCON problem are not known.

5. Polynomially solvable special cases

We have already remarked that the general problem of designing survivable networks is NP-hard; in fact, quite a number of special cases are known to be NP-hard. The intention of this section is to give an overview of those special cases where polynomial time solution procedures are known.

There are basically three ways to restrict the general model. One either considers special choices of node types, special cost functions or special classes of graphs. It turns out that some combinations of these restrictions lead to easy (but still interesting) problems.

5.1. Node type restrictions

Let us start by considering restrictions on the node types. $G = (V, E)$ may be any graph with costs $c_e \in \mathbf{R}$ for all $e \in E$.

If $r_v = 1$ for all $v \in V$ and $c_e > 0$ for all $e \in E$, the 1NCON and 1ECON problems for G and r are equivalent to finding a minimum cost spanning tree. This problem is well known to be solvable in polynomial time; see Chapter 12. If the

costs are not necessarily positive, we look for a minimum cost connected subgraph. This can be found by first choosing all edges of nonpositive cost, shrinking all resulting components to single nodes, and computing a minimum spanning tree in the resulting graph.

If two nodes have type 1, say nodes u and v , all other nodes have type 0, and if all costs are positive, then the 1ECON (and 1NCON) problem for G and r is nothing but the problem of computing a shortest $[u, v]$ -path; see Chapter 1 for polynomial time methods for that problem. If costs are general, we seek an edge set of minimum cost that contains a path from u to v . This can be solved in polynomial time by shrinking away the components induced by the nonpositive edges and then computing a shortest $[u, v]$ -path in the resulting graph.

The 'slightly' more general case, when $r_v \in \{0, 1\}$ for all $v \in V$, is the Steiner tree problem, which is NP-hard even if $c_e = 1$ for all $e \in E$. If there is a fixed number of nodes of type 0 or a fixed number of nodes of type 1, then the Steiner tree problem is polynomially solvable, see Lawler [1976].

The shortest path problem has an extension that is solvable in polynomial time. Namely, if two nodes have type k , say $r_u = r_v = k$, all others type 0, and if all costs are positive, then the k NCON problem asks for a collection of k node-disjoint $[u, v]$ -paths that are of minimum total cost and, similarly, the k ECON problem requires finding a minimum collection of k edge-disjoint $[u, v]$ -paths. Both problems can be solved with min-cost flow algorithms in polynomial time; see Chapter 1. As above, it is trivial to extend this to the case where costs are arbitrary.

We do not know of any other (nontrivial) case where special choices of node types lead to polynomial time solvability.

5.2. Cost restrictions

When edge costs are restricted to be in $\{1\}$ or $\{0, 1\}$, and the underlying graph is complete, two well-known problems of graph theory are obtained, namely, first, the problem of constructing graphs of certain connectivity properties having a minimum number of edges, and, second, the 'augmentation problem' of extending a given graph with as few as possible edges until it satisfies certain connectivity properties. We are going to survey some results for these problems.

Consider the following problem:

Problem 1. *Given a set V of nodes, and a requirement $r_{st} \geq 0$, for each pair $s, t \in V$, find a graph $G = (V, E)$ such that each pair s, t is at least r_{st} -edge connected in G and such that $|E|$ is as small as possible.*

Chou & Frank [1970] gave a polynomial-time algorithm to solve this problem when G may contain parallel edges, and when the edge-connectivity requirements are all at least 2. Frank & Chou [1970] solved a similar problem when no parallel edges but additional nodes are allowed in the construction. If neither parallel

edges nor further nodes are allowed, it is not known whether one can solve Problem 1 in polynomial time.

The node connectivity problem analogous to Problem 1 is open. We are only aware of a result of Harary [1962] who proved that, given n and k , the minimum number of edges in a k -connected graph on n nodes (without parallel edges) is $\lceil kn/2 \rceil$. Such a graph can be constructed easily.

To our knowledge these are the only solved cases with uniform edge costs.

The following problem often runs under the name *augmentation problem* in the graph theory literature.

Problem 2. *Given a graph, augment it by a minimum number of edges so that the new graph meets certain connectivity requirements.*

This type of problem was solved by Eswaran & Tarjan [1976] for 2-edge connected graphs. The k -edge connected graph augmentation problem was studied by Watanabe & Nakamura [1987], Ueno, Kajitani & Wada [1988], Cai & Sun [1989], and Naor, Gusfield & Martel [1990]. Frank [1992a] solved the general edge connectivity case. All these edge-connectivity augmentation algorithms allow the use of parallel edges. The solutions are algorithmic and can be found in polynomial time.

Frank [1992a] proved a nice min-max result for the minimum number of edges needed to augment a given graph G to satisfy given edge connectivity requirements r_{ij} . Let us define the *deficit* $\text{def}(A)$ of a node set A as

$$\text{def}(A) := \max_{u \in A, v \notin A} r_{uv} - |\delta_G(A)|.$$

The deficit of A is the smallest number of edges that have to be added to $\delta_G(A)$ in order to connect A sufficiently to all other nodes. Clearly, if several disjoint sets A_i ($i = 1, \dots, t$) have deficit $\text{def}(A_i)$, then a lower bound on the number of edges to be added is

$$\frac{1}{2} \sum_{i=1}^t \text{def}(A_i).$$

Frank [1992a] shows that, under certain assumptions, the best such lower bound is exactly the minimum number of edges needed in an augmentation.

Theorem 6. *Given a graph $G = (V, E)$ and edge connectivity requirements r_{st} for all pairs of nodes $s, t \in V$. Then the following holds:*

If some component of G with node set A has deficit at most 1 and all proper subsets of A have nonpositive deficit, then the minimum number of edges to be added to G is $\text{def}(A)$ plus the minimum number of edges to be added to $G - A$.

If no such component exists in G then the minimum number of edges to be added to G to satisfy the edge connectivity requirements, is

$$\max_{A_1, \dots, A_t \text{ disjoint}} \frac{1}{2} \sum_{i=1}^t \text{def}(A_i).$$

Frank's proof is constructive and results in a polynomial-time algorithm to create a minimum cost augmentation.

Augmentation to k -node connected graphs has been solved for $k = 2$ by Eswaran & Tarjan [1976], and for $k = 3$ by Hsu & Ramachandran [1991]. For $k = 4$, Hsu [1992] solved the problem, when the given graph is already 3-node connected. For general node connectivity requirements it is not even known whether the augmentation problem is NP-complete or not.

When directed graphs are considered, augmentation to k -edge connected digraphs is polynomially solvable, see Frank (1992a) (with parallel edges allowed), and so is augmentation to k -node connected digraphs, see Frank & Jordán [1993]. For general edge- or node-connectivity requirements, the augmentation problem in directed graphs was shown to be NP-complete by Frank [1992a].

5.3. Special classes of graphs

It is often the case that NP-hard problems become easy when restricted to graphs with special structural properties. In the case of the ECON and NCON problems, there are only very few (and relatively simple) classes of graphs known where some versions can be solved in polynomial time. Most notable is the class of series-parallel graphs.

Series-parallel graphs are created from a single edge by iterative application of the following two operations:

- addition of parallel edges, and
- subdivision of edges.

(For our purposes we note that all series-parallel graphs, except K_2 , are 2-connected.) *Outerplanar graphs* are a subclass of series-parallel graphs, namely those graphs that can be drawn in the plane as one cycle with noncrossing chords. *Halin graphs* are planar graphs that can be drawn in the plane as a tree without nodes of degree 2 plus one cycle connecting all leaves of the tree.

The Steiner tree problem can be solved in linear time on series-parallel graphs by a recursive algorithm. This was mentioned by Takamizawa, Nishizeki & Saito [1982], and stated explicitly by Wald & Colbourn [1983]. By a modification of this recursive algorithm, the 2NCON problem can also be solved, where node types 0, 1, and 2 are allowed.

Winter has developed linear-time algorithms for 2ECON and 2NCON problems with node types 0 and 2 on outerplanar, series-parallel, and Halin graphs, see Winter [1985a, b, 1986]. In his survey article, Winter [1987] mentioned that he also found linear-time algorithms for Halin graphs that solve the 3ECON and 3NCON problems with node types 0 and 3.

If there exist polynomial-time algorithms to solve edge- or node-connectivity problems on special classes of graphs then it is, in principle, possible to find a complete characterization by linear inequalities of the associated polytopes. Such characterizations are known for series-parallel graphs, and are listed in the following. Complete descriptions of Steiner tree polytopes and related polyhedra, using auxiliary variables, can be found in Prodon, Liebling & Gröflin

[1985], Goemans [1994a, b], Goemans & Myung [1993], and Margot, Prodon & Liebling [1994]. Using projection techniques one can obtain a complete characterization for the 1ECON polyhedron on series-parallel graphs without auxiliary variables, see also Goemans [1994b]. The projection technique and the inequalities generated by it are described in Section 8. A nonredundant description for the 1ECON polytope on series-parallel graphs is, however, yet unknown. Cornuéjols, Fonlupt & Naddef [1985] found a complete description of the dominant of the 2-edge connected subgraph polytope of series-parallel graphs. A complete description of the 2-edge-connected Steiner subgraph polytope on series-parallel subgraphs is given in Baiou & Mahjoub [1993]. For odd k , Chopra [1994] investigated the k -edge connected subgraphs of a given outerplanar graph, and found a complete description of the dominant of the associated polyhedron by the so-called *lifted outerplanar partition inequalities*. Chopra's result is as follows:

Theorem 7. *For outerplanar graphs $G = (V, E)$ and uniform node types $r \in \{k\}^V$, k odd, the dominant of the k ECON($G; r$) polytope (that is, k ECON($G; r$) + \mathbf{R}_+^E) is completely characterized by the inequalities:*

$$\frac{1}{2} \sum_{i=1}^p x(\delta(W_i)) \geq p \cdot \lceil k/2 \rceil - 1 \quad \text{for all partitions } \{W_1, \dots, W_p\} \text{ of } V_j,$$

$$x_e \geq 0 \quad \text{for all } e \text{ in } E.$$

This inequality can be lifted to an inequality valid and nonredundant for the dominant of k ECON(K_n, r) by computing the coefficients of the missing edges as the shortest-path value between their endpoints, and using as 'lengths' the coefficients on E).

6. Polyhedral results

Except for the results of Grötschel & Monma [1990] mentioned in Section 3, there is not much known about the polytope CON($G; r, k, d$) for general edge and node survivability requirements r, k and d . We will thus concentrate on the k NCON and k ECON problems that have been investigated in more depth and survey some of the known results. Particular attention has been paid to the low-connectivity case, that is, where $r \in \{0, 1, 2\}^E$. See Grötschel & Padberg [1985] and Pulleyblank [1989] for a general survey of polyhedral combinatorics and the basics of polyhedral theory.

Let us mention again the idea behind this approach and its goal. We consider an integer programming problem like (3) or (6). We want to turn such an integer program into a linear program and solve it using the (quite advanced) techniques of this area. To do this, we define a polytope associated with the problem by taking the convex hull of the feasible (integral) solutions of a program like (3) or (6). Let P be such a convex hull. We know from linear programming theory that, for

any objective function c , the linear program $\min\{c^T x \mid x \in P\}$ has an optimum vertex solution (if it has a solution). This vertex solution is, by definition, a feasible solution of the initial integer program and thus, by construction, an optimum solution of this program.

The difficulty with this approach is that $\max c^T x, x \in P$ is a linear program only ‘in principle’. To provide an instance to an LP-solver, we have to find a different description of P . The polytope P is defined as the convex hull of (usually many) points in \mathbf{R}^E , but we need a complete (linear) descriptions of P by means of linear equations or inequalities. The Weyl–Minkowski theorem tells us that both descriptions are in a sense equivalent, in fact, there are constructive procedures that compute one description of P from the other. However, these procedures are inherently exponential and nobody knows how to make effective use of them, in particular, for NP-hard problem classes. Moreover, there are results in complexity theory, see Papadimitriou & Yannakakis [1982], that indicate that it might be much harder to find a complete linear description of such a polytope P than to solve $\min c^T x, x \in P$.

At present, no effective general techniques are known for finding complete or ‘good partial’ descriptions of such a polytope or large classes of facets. There are a few basic techniques like the derivation of so-called Chvátal cuts (see Chvátal [1973]). But most of the work is a kind of ‘art’. Valid inequalities are derived from structural insights and the proofs that many of these inequalities define facets use technically complicated, ad-hoc arguments.

If large classes of facet-defining inequalities are found, one has to think about their algorithmic use. The standard technique is to employ such inequalities in the framework of a cutting plane algorithm. We will explain this in Section 7. It has turned out in the recent years that such efforts seem worthwhile. If one wants to find true optimum solutions or extremely good lower bounds, the methods of polyhedral combinatorics are the route to take.

6.1. Classes of valid inequalities

We will now give an overview of some of the results of Grötschel, Monma & Stoer [1992a–c] and Stoer [1992] concerning classes of valid inequalities for the k ECON and k NCON problems. We will motivate these inequalities and mention how they arise.

As before, we consider a loopless graph $G = (V, E)$, and in the k ECON case possibly with multiple edges. We assume that for each node $v \in V$ a nonnegative integer r_v , its node type, is given, that $k = \max\{r_v \mid v \in V\}$ and that at least two nodes are of type k . Recall that $r(W) = \max\{r_v \mid v \in W\}$ is called the node type of W .

We start out by repeating those classes we have already introduced in Section 3. Clearly, the *trivial inequalities*

$$0 \leq x_e \leq 1 \quad \text{for all } e \in E \quad (7)$$

are valid for k ECON($G; r$) and k NCON($G; r$) since problem (6) is a 0/1-

optimization problem. The *cut inequalities*

$$x(\delta(W)) \geq \text{con}(W) \quad \text{for all } W \subseteq V, \emptyset \neq W \neq V, \tag{8}$$

where $\text{con}(W)$ is given by (1), or equivalently by $\min\{r(W), r(V \setminus W)\}$, are valid for $k\text{ECON}(G; r)$ and $k\text{NCON}(G; r)$, since the survivable network to be designed has to contain at least $\text{con}(W)$ edge-disjoint paths that connect nodes in W to nodes in $V \setminus W$. (Recall that $r_{st} = \min\{r_s, r_t\}$, $s, t \in V$.) In the node connectivity case we require that upon deletion of any set Z of nodes there has to be, for all pairs $s, t \in V \setminus Z$, at least $r_{st} - |Z|$ more paths connecting s and t in the remaining graph. This requirement leads to the *node cut inequalities*

$$x(\delta_{G-Z}(W)) \geq \text{con}(W) - |Z| \tag{9}$$

for all pairs $s, t \in V, s \neq t$ and
 for all $\emptyset \neq Z \subseteq V \setminus \{s, t\}$ with $|Z| \leq r_{st} - 1$
 and for all $W \subseteq V \setminus Z$ with $s \in W, t \notin W$.

These inequalities are valid for $k\text{NCON}(G; r)$ but — of course — not for $k\text{ECON}(G; r)$.

How does one find further classes of valid inequalities? One approach is to infer inequalities from structural investigations. For instance, the cut inequalities ensure that every cut separating two nodes contains at least r_{st} edges. These correspond to partitioning the node set into two parts and guaranteeing that there are enough edges linking them. We can generalize this idea as follows. Let us call a system W_1, \dots, W_p of nonempty subsets of V with $W_i \cap W_j = \emptyset$ for $1 \leq i < j \leq p$, and $W_1 \cup \dots \cup W_p = V$ a *partition* of V and let us call

$$\begin{aligned} \delta(W_1, \dots, W_p) &:= \\ &:= \{uv \in E \mid \exists i, j, 1 \leq i, j \leq p, i \neq j \text{ with } u \in W_i, v \in W_j\} \end{aligned}$$

a *multicut* or *p-cut* (if we want to specify the number p of *shores* W_1, \dots, W_p of the multicut). Depending on the numbers $\text{con}(W_1), \dots, \text{con}(W_p)$, any survivable network (V, F) will have to contain at least a certain number of edges of the multicut $\delta(W_1, \dots, W_p)$. For every partition it is possible to compute a lower bound of this number, and thus to derive a valid inequality for every node partition (resp. multicut). This goes as follows.

Suppose W_1, \dots, W_p is a partition of V such that $\text{con}(W_i) \geq 1$ for $i = 1, \dots, p$. Let $I_1 := \{i \in \{1, \dots, p\} \mid \text{con}(W_i) = 1\}$, and $I_2 := \{i \in \{1, \dots, p\} \mid \text{con}(W_i) \geq 2\}$. Then the *partition inequality* (or multicut inequality) induced by W_1, \dots, W_p is defined as

$$\begin{aligned} x(\delta(W_1, \dots, W_p)) &= \\ &= \frac{1}{2} \sum_{i=1}^p x(\delta(W_i)) \geq \begin{cases} \left\lceil \frac{1}{2} \sum_{i \in I_2} \text{con}(W_i) \right\rceil + |I_1| & \text{if } I_2 \neq \emptyset, \\ p - 1 & \text{if } I_2 = \emptyset. \end{cases} \end{aligned} \tag{10}$$

Every partition inequality is valid for $k\text{ECON}(G; r)$ and thus for $k\text{NCON}(G; r)$.

Just as the cut inequalities (8) can be generalized as outlined above to partition inequalities (10), the node cut inequalities (9) can be generalized to a class of inequalities that we will call node partition inequalities, as follows.

Let $Z \subseteq V$ be some node set with $|Z| \geq 1$. If we delete Z from G then the resulting graph must contain an $[s, t]$ -path for every pair of nodes s, t of type larger than $|Z|$. In other words, if W_1, \dots, W_p is a partition of $V \setminus Z$ into node sets with $r(W_i) \geq |Z| + 1$ then the graph G' obtained by deleting Z and contracting W_1, W_2, \dots, W_p must be connected. This observation gives the following class of *node partition inequalities* valid for $kNCON(G; r)$, but not for $kECON(G; r)$:

$$\frac{1}{2} \sum_{i=1}^p x(\delta_{G-Z}(W_i)) \geq p - 1 \quad (11)$$

for every node set $Z \subseteq V$, $|Z| \geq 1$ and
every partition W_1, \dots, W_p of $V \setminus Z$
such that $r(W_i) \geq |Z| + 1$, $i = 1, \dots, p$.

If $r(W_i) \geq |Z| + 2$ for at least two node sets in the partition, then the right-hand side of the node partition inequality can be increased. This leads to further generalizations of the classes (10) and (11), but their description is quite technical and complicated, see Stoer [1992]. So we do not discuss them here.

We now mention another approach to finding new classes of valid inequalities.

The idea here is to relax the problem in question by finding a (hopefully easier) further combinatorial optimization problem such that every solution of the given problem is feasible for the new problem. One can then study the polytope associated with the new combinatorial optimization problem. If the relaxation is carefully chosen — and one is lucky — some valid inequalities for the relaxed polytope turn out to be facet-defining for the polytope one wants to consider. These inequalities are trivially valid.

In our case, a relaxation that is self-suggesting is the so-called r -cover problem. This ties the survivability problem to matching theory and, in fact, one can make good use of the results of this theory for the survivability problem.

The survivability requirements imply that if $v \in V$ is a node of type r_v , then v has degree at least r_v for any feasible solution of the $kECON$ problem. Thus, if we can find an edge set of minimum cost such that each node has degree at least r_v (we call such a set an r -cover), we obtain a lower bound for the optimum value of the $kECON$ problem. Clearly, such an edge set can be found by solving the integer linear program

$$\begin{aligned} \min \quad & c^T x \\ \text{(i)} \quad & x(\delta(v)) \geq r_v \text{ for all } v \in V, \\ \text{(ii)} \quad & 0 \leq x_e \leq 1 \text{ for all } e \in E, \text{ and} \\ \text{(iii)} \quad & x_e \text{ integer for all } e \in E, \end{aligned} \quad (12)$$

which is obtained from (3i), (3iii) and (3iv) by considering only sets of cardinality one in (3i). The inequalities (12i) are called *degree constraints*. This integer program can be turned into a linear program, i.e., the integrality constraints (12iii)

are replaced by a system of linear inequalities, using Edmonds' polyhedral results on b -matching, see Edmonds [1965]. Edmonds proved that, for any vector $b \in Z_+^V$, the vertices of the polyhedron defined by

$$\begin{aligned} \text{(i)} \quad & y(\delta(v)) \leq b_v && \text{for all } v \in V, \\ \text{(ii)} \quad & y(E(H)) + y(\bar{T}) \leq \left\lfloor \frac{1}{2} \sum_{v \in H} (b_v + |\bar{T}|) \right\rfloor && \text{for all } H \subseteq V \\ & && \text{and all } \bar{T} \subseteq \delta(H), \text{ and} \\ \text{(iii)} \quad & 0 \leq y_e \leq 1 && \text{for all } e \in E \end{aligned} \quad (13)$$

are precisely the incidence vectors of all (1-capacitated) b -matchings of G , i.e., of edge sets M such that no node $v \in V$ is contained in more than b_v edges of M . For the case $b_v := |\delta(v)| - r_v$, the b -matchings M are nothing but the complements $M = E \setminus F$ of r -covers F of G . Using the transformation $x := 1 - y$ and $T := \delta(H) \setminus \bar{T}$ we obtain the system

$$\begin{aligned} \text{(i)} \quad & x(\delta(v)) \geq r_v && \text{for all } v \in V, \\ \text{(ii)} \quad & x(E(H)) + x(\delta(H) \setminus T) \geq \left\lfloor \frac{1}{2} \sum_{v \in H} (r_v - |T|) \right\rfloor && \text{for all } H \subseteq V \\ & && \text{and all } T \subseteq \delta(H), \text{ and} \\ \text{(iii)} \quad & 0 \leq x_e \leq 1 && \text{for all } e \in E. \end{aligned} \quad (14)$$

(14) gives a complete description of the convex hull of the incidence vectors of all r -covers of G . We call the inequalities (14ii) r -cover inequalities. Since every solution of the k ECON problem for G and r is an r -cover, all inequalities (14ii) are valid for k ECON($G; r$). It is a trivial matter to observe that those inequalities (14ii) where $\sum_{v \in H} r_v - |T|$ is even are redundant. For the case $r_v = 2$ for all $v \in V$, Mahjoub [1994] described the class of r -cover inequalities, which he calls odd wheel inequalities.

Based on these observations one can extend inequalities (14ii) to more general classes of inequalities valid for k ECON($G; r$) (but possibly not valid for the r -cover polytope). We present here one such generalization.

Let H be a subset of V called the *handle*, and $T \subseteq \delta(H)$ with $|T|$ odd and $|T| \geq 3$. For each $e \in T$, let T_e denote the set of the two end nodes of e . The sets $T_e, e \in T$, are called *teeth*. Let H_1, \dots, H_p be a partition of H into nonempty pairwise disjoint subsets such that $r(H_i) \geq 1$ for $i = 1, \dots, p$, and $|H_i \cap T_e| \leq r(H_i) - 1$ for all $i \in \{1, \dots, p\}$ and all $e \in T$. Let $I_1 := \{i \in \{1, \dots, p\} \mid r(H_i) = 1\}$ and $I_2 = \{i \in \{1, \dots, p\} \mid r(H_i) \geq 2\}$. We call

$$x(E(H)) - \sum_{i=1}^p x(E(H_i)) + x(\delta(H) \setminus T) \geq \left\lfloor \frac{1}{2} \sum_{i \in I_2} (r(H_i) - |T|) \right\rfloor + |I_1| \quad (15)$$

the *lifted r -cover inequality* (induced by H_1, \dots, H_p, T). All inequalities of type (15) are valid for k ECON($G; r$).

The names 'handle' and 'teeth' used above derive from the observation that there is some relationship of these types of inequalities with the 2-matching, comb and clique tree inequalities for the symmetric traveling salesman polytope; see Chapter 3. In fact, comb inequalities for the traveling salesman problem can

be transformed in various ways to facet-defining inequalities for 2ECON and k NCON polyhedra, as mentioned in Grötschel, Monma & Stoer [1992a], Boyd & Hao [1993], and Stoer [1992].

Another technique for finding further classes of valid and facet-defining inequalities will be mentioned in Section 6.3.

To develop successful cutting plane algorithms, it is not enough to know some inequalities valid for the polytope over which one wants to optimize. The classes of inequalities should contain large numbers of facets of the polytope. Ideally, one would like to use classes of facet-defining inequalities only.

In our case, it turned out to be extremely complicated to give (checkable) necessary and sufficient conditions for an inequality in one of the classes described above to define a facet of k NCON($G; r$) or k ECON($G; r$). Lots of technicalities creep in, when general graphs G , as opposed to complete graphs, are considered. Nevertheless, it could be shown that large subsets of these classes are facet-defining also for the relatively sparse graphs that come from the applications, see Figures 4 and 7 for examples. These results provide a theoretical justification for the use of these inequalities in a cutting plane algorithm. Details about facet results for the inequalities described above can be found in Grötschel & Monma [1990], Grötschel, Monma & Stoer [1992a–c], Stoer [1992].

6.2. Separation

Note that — except for the trivial inequalities — all classes of valid inequalities for the k ECON and k NCON problem described in Section 6.1 contain a number of inequalities that is exponential in the number of nodes of the given graph. So it is impossible to input these inequalities into an LP-solver. But there is an alternative approach. Instead of solving an LP with all inequalities, we solve one with a few ‘carefully selected’ inequalities and we generate new inequalities as we need them. This approach is called a cutting plane algorithm and works as follows.

We start with an initial linear program. In our case, it consists of the linear program (12) without the integrality constraints (12iii). We solve this LP. If the optimum solution y is feasible for the k ECON or k NCON problem, then we are done. Otherwise we have to find some inequalities that are valid for k ECON($G; r$) or k NCON($G; r$) but are violated by y . We add these inequalities to the current LP and repeat.

The main difficulty of this approach is in efficiently generating violated inequalities. We state this task formally.

Separation Problem (for a class C of inequalities). *Given a vector y decide whether y satisfies all inequalities in C and, if not, output an inequality violated by y .*

A trivial way to solve Problem 3 is to substitute y into each of the inequalities in C and check whether one of the inequalities is violated. But in our case this is too time consuming since C is of size exponential in $|V|$. Note that all the classes C

described before have an implicit description by means of a formula with which all inequalities can be generated. It thus may happen that algorithms can be designed that check violation much more efficiently than the trivial substitution process. We call an algorithm that solves Problem 3 an (exact) *separation algorithm* for C , and we say that it runs in polynomial time if its running time is bounded by a polynomial in $|V|$ and the encoding length of y .

A deep result of the theory of linear programming, see Grötschel, Lovász & Schrijver [1988], states (roughly) that a linear program over a class C of inequalities can be solved in polynomial time if and only if the separation problem for C can be solved in polynomial time. Being able to solve the separation problem thus has considerable theoretical consequences.

This result makes use of the ellipsoid method and does not imply the existence of a 'practically efficient' algorithm. However, the combination of separation algorithms with other LP solvers (like the simplex algorithms) can result in quite successful cutting plane algorithms; see Section 7.

Our task now is to find out whether reasonable separation algorithms can be designed for any of the classes (8), (9), (10), (11), (14ii), and (15).

There is some good and some bad news. The good news is that for the cut inequalities (8), the node cut inequalities (9) and the r -cover inequalities (14ii), exact separation algorithms are known that run in polynomial time; see Grötschel, Monma & Stoer [1992c].

When C is the class of cut inequalities and y is a nonnegative vector, separation can be solved by any algorithm determining a cut $\delta(W)$ of minimum capacity $y(\delta(W))$ in a graph. Fast min-cut algorithms are described in Hao & Orlin [1992] and Nagamochi & Ibaraki [1992]. Both algorithms do not need more than $O(|V|^3)$ time. The so-called Gomory–Hu tree storing one minimum (s, t) -cut for each pair of nodes s, t in a tree structure can be computed in $O(|V|^4)$ time, see Gomory & Hu [1961].

When C is the class of cut and node cut inequalities (8) and (9), the separation problem can be reduced to a sequence of minimum (s, t) -cut computations in a directed graph. This polynomial-time method is described in Grötschel, Monma & Stoer [1992c].

The polynomial-time exact separation algorithm for the r -cover inequalities is based on the Padberg–Rao procedure for solving the separation problem for the capacitated b -matching inequalities, see Padberg & Rao [1982]. The 'trick' is to reverse the transformation from the b -matching to the r -cover problem described in (13) and (14) and call the Padberg–Rao algorithm. It is easy to see that y satisfies all r -cover inequalities (14ii) if and only if its transformation satisfies all b -matching inequalities (13ii). The Padberg–Rao procedure is quite complicated to describe, so we do not discuss it here.

The bad news is that it was shown in Grötschel, Monma & Stoer [1992b] that the separation problems for partition inequalities (10), node partition inequalities (11) and lifted r -cover inequalities (15) are NP-hard. (A certain generalization of partition inequalities for k ECON problems with $k \leq 2$ is, however, polynomial-time separable, see Section 8).

Thus, in these cases we have to revert to *separation heuristics*, i.e., fast procedures that check whether they can find an inequality in the given class that is violated by y , but which are not guaranteed to find one even if one exists. We discuss separation heuristics in more detail in Section 7.

6.3. Complete descriptions of small cases

For the investigation of combinatorially defined polyhedra, it is often useful to study small cases first. A detailed examination of such examples provides insight into the relationship between such polyhedra and gives rise to conjectures about general properties of these polyhedra. Quite frequently a certain inequality is found, by numerical computation, to define a facet of a k ECON or k NCON polytope of small dimension. Afterwards it is often possible to come up with a class (or several classes) of inequalities that generalize the given one and to prove that many inequalities of these classes define facets of combinatorial polytopes of any dimension.

By means of a computer program, we have computed complete descriptions of all k ECON and k NCON polytopes of small dimensions. To give a glimpse of these numerically obtained results, we report here the complete descriptions of all 2ECON and all 2NCON polytopes of the complete graphs on five vertices K_5 . More information about small k ECON polytopes can be found in Stoer [1992].

6.3.1. The 2ECON polytope for K_5

Let us begin with the polytopes $2ECON(K_5; r)$ where $r = (r_1, \dots, r_5)$ is the vector of node types. The node types r_i have value 0, 1 or 2, and by assumption, at least two nodes are of highest type 2. Clearly, we can suppose that $r_i \geq r_{i+1}$ for $i = 1, \dots, 4$. These assumptions result in ten node type vectors to be considered. It is obvious that, if a node type vector r componentwise dominates a vector r' (i.e., $r_i \geq r'_i$ for all i), then $2ECON(K_5; r)$ is contained in $2ECON(K_5; r')$.

Figure 1 provides a comprehensive summary of our findings. In this figure, a polytope $2ECON(K_5; r)$ is depicted by its node type vector $r = (r_1, \dots, r_5)$. A line linking two such vectors indicates that the polytope at the lower end of the line directly contains the polytope at the upper end of the line and that no other 2ECON polytope is 'in between'. For example, the polytope $2ECON(K_5; (2, 2, 2, 1, 0))$ is directly contained in $2ECON(K_5; (2, 2, 1, 1, 0))$ and $2ECON(K_5; (2, 2, 2, 0, 0))$, and it contains directly $2ECON(K_5; (2, 2, 2, 1, 1))$ and $2ECON(K_5; (2, 2, 2, 2, 0))$. Next to the right or left of a node type vector r , a box indicates which inequalities are needed to define the corresponding 2ECON $(K_5; r)$ polytope completely and nonredundantly. The notation is as follows:

- The type of inequality appears in the first column.
- 'p' stands for 'partition inequality', see (10),
- 'cut' stands for 'cut constraint', see (8),
- 'rc' stands for 'lifted r -cover inequality', see (15),
- 'd' stands for 'degree constraint', see (12i),
- 'rc+1', 'I1', and 'I2', stand for new types of inequalities explained later.

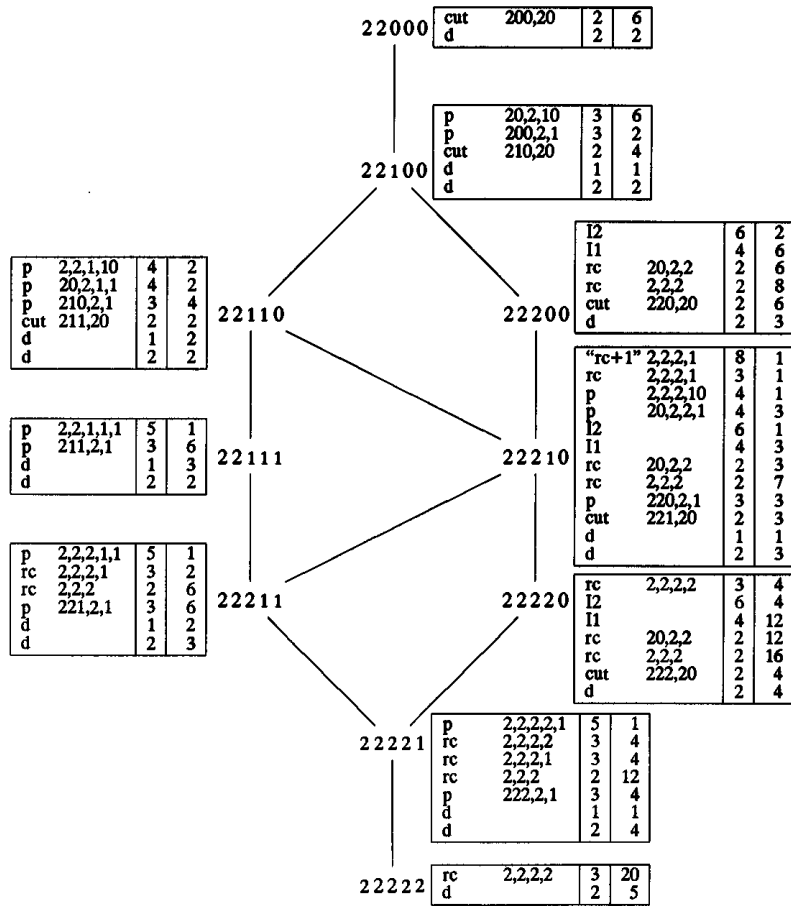


Fig. 1. 2ECON($K_5; r$) polyhedra.

- The next column lists, for each partition of the handle ('rc') or the whole node set ('p'), the node types in each node set. The different sets are separated by commas.

- For instance, 'p 200,2,1' stands for a partition inequality induced by a partition of V , whose first node set contains one node of type 2 and two nodes of type 0, whose second set contains exactly one node of type 2, and whose last set contains exactly one node of type 1.
- 'rc 20,2,2' stands for a lifted r -cover inequality induced by a handle that is partitioned into three node sets, the first one containing two nodes, a node of type 2 and a node of type 0, the second node set containing exactly one node of type 2, and the third node set containing exactly one node of type 2; the number of teeth can be computed with the help of the right-hand side.

- The next column gives the right-hand side.
- The last column contains the number of inequalities of the given type.

We do not list the trivial inequalities $0 \leq x_e \leq 1$, since they always define facets of the considered polytopes, see Theorem 2.

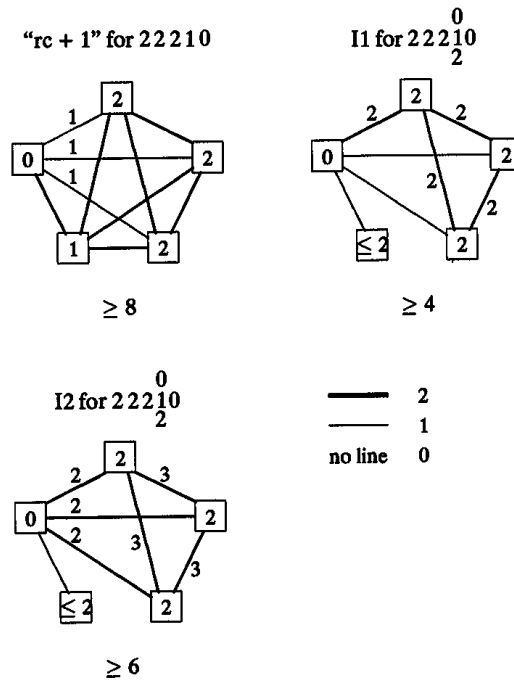


Fig. 2. Inequalities for 2ECON($K_5; r$).

The inequalities denoted by 'rc+1', 'I1', and 'I2' in 2ECON($K_5; (2, 2, 2, 0, 0)$), and 2ECON($K_5; (2, 2, 2, 1, 0)$), are depicted in Figure 2. All except I2 have coefficients in $\{0, 1, 2\}$. The coefficients of inequality I2 in Figure 2 take values in $\{0, 1, 2, 3\}$. Edges of coefficient 0 are not drawn, edges of coefficient 1 are drawn by thin lines, and edges of coefficient 2 are drawn by bold lines. To make this distinction somewhat clearer, we additionally display the coefficients of all thin or of all bold lines.

This numerical study of 2ECON polytopes of K_5 reveals that degree, cut, partition, and lifted r -cover inequalities play a major role in describing the polytopes completely and nonredundantly. But at the same time we found three new classes of (complicated) inequalities.

6.3.2. The 2NCON polytope for K_5

We now turn our attention to the 2NCON polytopes of the complete graph K_5 . It turned out that only two further classes of inequalities are needed to describe all polytopes 2NCON ($K_5; r$) completely and nonredundantly. These are the classes of node cut and node partition inequalities (9) and (11). Figure 3 displays the complete descriptions of the 2NCON polytopes for K_5 in the same way as Figure 1 does for the 2ECON polytopes. The (new) entries for the node cut and node partition inequalities read as follows.

'ncut 20,20' denotes a node cut inequality $x(\delta_{G-z}(W))$, where both W and $V \setminus (W \cup \{z\})$ contains a node of type 2 and a node of type 0, and the '.' in 'ncut 2.,2.' represents a node of any type;

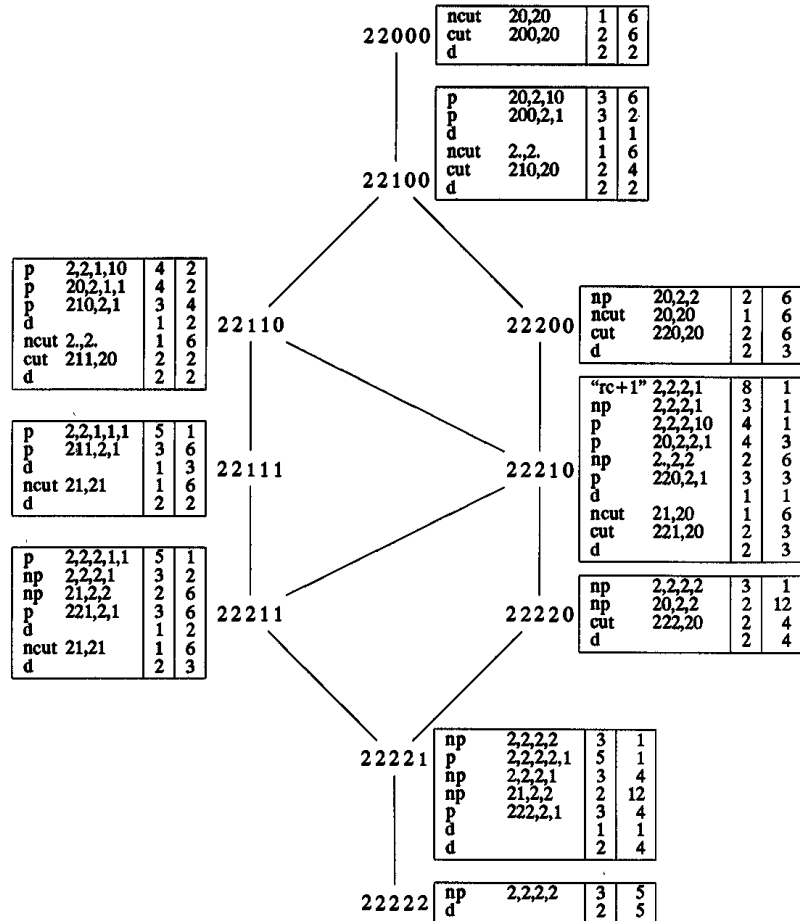


Fig. 3. 2NCON($K_5; r$) polyhedra.

and 'np 20,2,2' denotes a node partition inequality induced by a partition of $V \setminus \{z\}$, where z is some node in V , and the first shore of the partition consists of a node of type 2 and a node of type 0, the second shore consists of a node of type 2, and the third shore consists of a node of type 2.

This concludes our examination of polyhedra for small instances of 2NCON and 2ECON problems. For further such results, see Stoer [1992].

7. Computational results

For applied mathematicians, the ultimate test of the power of a theory is its success in helping solve the practical problems for which it was developed. In our case, this means that we have to determine whether the polyhedral theory for the survivable network design problem can be used in the framework of a cutting plane algorithm to solve the design problems of the sizes arising in practice. The results reported in Grötschel, Monma & Stoer [1992b] show that

the design problems for Local Access Transport Area (LATA) networks arising at Bell Communications Research (Bellcore) can be solved to optimality quite easily. There is good reason to hope that other network design problems of this type can also be attacked successfully with this approach. Moreover, the known heuristics also seem to work quite well, at least for the low connectivity case.

7.1. Outline of the cutting plane algorithm

We have already mentioned in Section 6.2 how a cutting plane approach for our problem works. Let us repeat this process here a little more formally.

We assume that a graph $G = (V, E)$ with edge cost $c_e \in \mathbf{R}$ for all $e \in E$ and node types $r_v \in \mathbf{Z}_+$ for all $v \in V$ is given. Let $k := \max\{r_v \mid v \in V\}$. We want to solve either the k NCON or the k ECON problem for G and r and the given cost function, i.e., we want to solve

$$\min_{x \in k\text{NCON}(G;r)} c^T x \quad \text{or} \quad \min_{x \in k\text{ECON}(G;r)} c^T x.$$

We do this by solving a sequence of linear programming relaxations that are based on the results we have described in Section 6. The initial LP (in both cases) consists of the degree constraints and trivial inequalities, i.e.,

$$\begin{aligned} \min \quad & c^T x \\ & x(\delta(v)) \geq r_v \quad \text{for all } v \in V; \\ & 0 \leq x_e \leq 1 \quad \text{for all } e \in E. \end{aligned} \tag{16}$$

If a current LP-relaxation has been solved and y is a basic optimum solution, we check whether y is in k NCON(G ; r) or k ECON(G ; r). If it is, the problem is already solved to optimality. Otherwise we call our separation routines and try to find inequalities in one of the classes described in Section 6 that are violated by y . If one of the separation algorithms is successful, we add the inequalities found to the current LP and repeat the procedure. If no violated inequality can be identified, there are two options. We may simply stop and report a lower bound or we may resort to an enumerative procedure like branch and bound.

The best option to select in the case that the cutting plane phase has not produced an optimum solution of the k ECON or k NCON problem depends on the demands of practice. Before starting the cutting plane algorithm, one usually runs some heuristics to produce (hopefully) good feasible solutions, see Section 4 for a description of such heuristics. By comparing the lower bound L from the cutting plane algorithm with the upper bound U from the heuristics, one can easily get an idea about the quality of the solutions. If the percentage deviation of these values, usually taken as $100 * (U - L) / L$, is less than some threshold, say 5%, and if the cost data were somewhat fuzzy anyway, the best heuristic solution might simply be considered as an appropriate good solution of the practical problem. If the data are precise and an optimum solution is needed, branch and bound has a good chance to terminate in a 'reasonable' amount of time.

If, however, the deviation of U and L is large, no simple advice can be given. Either the heuristic or the cutting plane algorithm or both may have produced poor bounds. Further research is usually necessary to determine the reasons for failure, to detect special structures of the problem in question that may have resulted in traps for the heuristics or poor performance of the cutting plane algorithms. If such structural insight can be obtained, one has a good chance to improve the result, at least for the current problem instance.

The last case is definitely unsatisfactory, but since we are dealing with hard problems, we have to expect such behavior every now and then. For our real world applications we can report that, for the LATA network design problems, the cutting plane algorithm always found an optimum integral solution, except in three cases that were easily solved by branch and bound or manual interaction. Random problems with low and high connectivity requirements and random cost structure were solved to optimality extremely quickly. But there is one large scale practical problem with 494 nodes, and 1096 edges and highly structured topology and costs, where we ran into considerable difficulties. Using special purpose separation heuristics etc., we were finally able to solve two versions of the problem to optimality, with quite some effort however.

7.2. Implementation details

We will not discuss implementation details of the heuristics outlined in Section 4; see Monma & Shallcross [1989] for details. This is quite straightforward, although of course, the use of good data structures and search strategies may result in considerable running time improvements.

We will focus here on implementation issues of the cutting plane algorithm, a number of which are vital for obtaining satisfactory running time performances. Before starting the cutting plane algorithm we try to reduce the size of the problem instance by decomposing it. In fact, the practical problems we solved have rather sparse graphs of possible direct links, and the survivability requirements often force certain edges to be present in every feasible solution. Such edges can be fixed and removed from the problem by appropriately changing certain node types. This removal may break the original problem into several smaller ones that can be solved independently.

There are further ways of decomposing a problem into independent subproblems like decomposing on articulation nodes, on cut sets of size two, and on articulation sets of size two. In each of these cases one can perform the decomposition or determine that no such decomposition is possible, using polynomial time methods like depth-first search or connectivity algorithms. All of this is quite easy, though a precise description of the necessary transformations would require considerable space. Details can be found in Grötschel, Monma & Stoer [1992b] and Stoer [1992].

The main purpose of this decomposition step is to speed up the computation by getting rid of some trivial special cases that the cutting plane algorithm does not need to check any more, and by reducing the sizes of the problems to be solved.

At the end of this preprocessing phase we have decomposed the original problem into a list of subproblems for which we call the cutting plane algorithm. The optimal solution of the original problem can then be composed from the optimal solutions of the subproblems in a straightforward manner.

An issue of particular importance is the implementation of the separation algorithms. Good heuristics have to be found that 'solve' the separation problems for those classes that are not known to be separable in polynomial time. And even if polynomial exact separation routines are known, separation heuristics may help considerably to speed up the overall program. Further problems are to determine the order in which the heuristic and exact separation routines are to be called, when to stop the separation process in case many violated inequalities have been found, which cutting planes to add and which to eliminate from the current LP. These issues cannot be decided by 'theoretical insight' only. Practical experience on many examples is necessary to come up with recipes that result in satisfactory overall performance of such an algorithm. We outline some of the techniques used in the sequel.

Let G be a graph with node types $r \in \mathbf{Z}_+^V$, and let y be some point in \mathbf{R}^E with $0 \leq y_e \leq 1$. Our aim is to find a partition (resp. node partition, lifted r -cover) inequality violated by this point. By the NP-completeness results mentioned in Section 6.2 it seems hopeless to find an efficient exact algorithm for the separation of these inequalities; therefore we have to use heuristics. Nevertheless, it is possible to solve the separation problem for a certain subclass of these inequalities, namely cut constraints (resp., node cut and r -cover constraints) in polynomial time. So in our heuristics we often use 'almost violated' inequalities of these subclasses and transform them into violated inequalities of the larger class. Here an almost violated inequality is an inequality $a^T x \geq b$ with $a^T y \leq b + \alpha$ for some 'small' parameter α (we used $\alpha = 0.5$). $a^T x \geq b$ is a violated inequality, if $a^T y < b$.

The heuristic that we applied has the following general form:

Heuristic for finding violated partition inequalities

(1) *Shrink all or some edges $e \in E$ with the property that any violated partition inequality using this edge can be transformed into some at-least-as-violated partition inequality not using this edge. ('Using e ' means; e has coefficient 1 in the partition inequality.)*

(2) *Find some violated or almost violated cut constraints in the resulting graph.*

(3) *Attempt to modify these cut constraints into violated partition inequalities.*

Exactly the same approach is used for separating node partition (11) and lifted r -cover inequalities (15), except that we have to use other shrinking criteria and, in Step 2, plug in the appropriate subroutine for separating node cut (8), resp., r -cover constraints (14ii).

Shrinking is important for reducing graph sizes before applying a min-cut algorithm and for increasing sizes of shores. If we are looking for related partition inequalities, we test whether edge $e = uv$ satisfies one of the following shrinking criteria.

Shrinking criteria

- (1) $y_e \geq q := \max\{r_w : w \in V\}$.
- (2) $y_e \geq r_v$ and $y_e \geq y(\delta(v)) - y_e$.
- (3) $y_e \geq \max\{y(\delta(v)) - y_e, y(\delta(u)) - y_e\}$ and there is a node $w \notin \{u, v\}$ with $r_w \geq \max\{r_u, r_v\}$.

If these criteria are satisfied for edge e , we shrink it by identifying u and v , giving type $\text{con}(\{u, v\})$ to the new node, and identifying parallel edges by adding their y -values.

It can be shown, that if cases (1) or (2) apply, then any violated partition inequality using e can be transformed into some at-least-as-violated partition inequality not using e . In case (3), edge e has the same property with respect to cut inequalities. Similar shrinking criteria can be found for node partition inequalities and lifted r -cover inequalities.

In the reduced graph G' we now find violated or almost violated cut constraints (resp., node partition and r -cover constraints) using the Gomory–Hu algorithm (or, for r -cover constraints the Padberg–Rao algorithm). These inequalities, defined for G' , are transformed back into the original graph G . For instance, a cut inequality in G' , $x(\delta_{G'}(W')) \geq r(W')$ is first transformed into a cut inequality $x(\delta_G(W)) \geq r(W)$ in G by blowing up all shrunk nodes in W' . This provides the enlarged node set W . Secondly, this cut inequality is transformed into a (hopefully) violated partition inequality by splitting W or $V \setminus W$ into smaller disjoint node sets W_1, \dots, W_p .

We also check whether the given cut inequality satisfies some simple necessary criteria for defining a facet of $k\text{ECON}(G; r)$ (or $k\text{NCON}(G; r)$). If this is not so, it can usually be transformed into a partition inequality that defines a higher-dimensional face of the respective polyhedron. A similar approach is taken for node partition and lifted r -cover inequalities. More details can be found in Grötschel, Monma & Stoer [1992b] and in Stoer [1992]. Typically, in the first few iterations of the cutting plane algorithm, the fractional solution y consists of several disconnected components. So, y violates many cut and partition inequalities but usually no lifted r -cover inequalities. We start to separate lifted r -cover inequalities only after the number of violated partition inequalities found drops below a certain threshold. Node partition inequalities are used only after all other separation algorithms failed in finding more than a certain number of inequalities.

To keep the number of LP constraints small, all inequalities in the current LP with non-zero slack are eliminated. But since all inequalities ever found by the separation algorithms are stored in some external pool they can be added again if violated at some later point.

7.3. Computational results for low-connectivity problems

In this section we describe the computational results based on the practical heuristics described in Section 4 and the cutting plane approach described earlier

Table 1
Data for LATA problems

Problem	Original graphs					Reduced graphs				
	0	1	2	Nodes	Edges	0	1	2	Nodes	Edges
LATADMA	0	12	24	36	65/0	0	6	15	21	46/4
LATA1	8	65	14	77	112/0	0	10	14	24	48/2
LATA5S	0	31	8	39	71/0	0	15	8	23	50/0
LATA5L	0	36	10	46	98/0	0	20	9	29	77/1
LATADSF	0	108	8	116	173/40	0	28	11	39	86/26
LATADS	0	108	8	116	173/0	0	28	11	39	86/3
LATADL	0	84	32	116	173/0	0	11	28	39	86/6

in this section. We consider the low-connectivity case with node types in $\{0, 1, 2\}$ here and the high connectivity case in the next section. All running times reported are for a SUN 4/50 IPX workstation (a 28.5 MIPS machine). The LP-solver used is a research version of the CPLEX-code provided to us by Bixby [1992]. This is a very fast implementation of the simplex algorithm.

To test our code, network designers at Bellcore provided the data (nodes, possible direct links, costs for establishing a link) of seven real LATA networks that were considered typical for this type of application. The sizes ranged from 36 nodes and 65 edges to 116 nodes and 173 edges; see Table 1. The problem instances LATADL, LATADS, and LATADSF are defined on the same graph. The edges have the same costs in each case, but the node types vary. Moreover, in LATADSF, 40 edges were required to be in the solution. (The purpose was to check how much the cost would increase if these edges had to be used, a typical situation in practice, where alternative solutions are investigated by requiring the use of certain direct links.)

Table 1 provides information about the problems. Column 1 contains the problem names. For the original graphs, columns 2, 3, and 4 contain the numbers of nodes of type 0, 1, and 2, respectively; column 5 lists the total number of nodes, column 6 the number of edges and the number of edges required to be in any solution (the forced edges). All graphs were analysed by our preprocessing procedures described in Section 7.2. Preprocessing was very successful. In fact, in every case, the decomposition and fixing techniques ended up with a single, much smaller graph obtained from the original graph by splitting off side branches consisting of nodes of type 1, replacing paths where all interior nodes are of degree 2, by a single edge, etc. The data of the resulting reduced graphs are listed in columns 6, . . . , 10 of Table 1.

To give a visual impression of the problem topologies and the reductions achieved, we show in Figure 4 a picture of the original graph of the LATADL problem (with 32 nodes of type 2 and 84 nodes of type 1) and in Figure 5 a picture of the reduced graph (with 39 nodes and 86 edges) after preprocessing. The nodes of type 2 are displayed by squares, and the nodes of type 1 are displayed by circles. The 6 forced edges that have to be in any feasible solution are drawn bold.

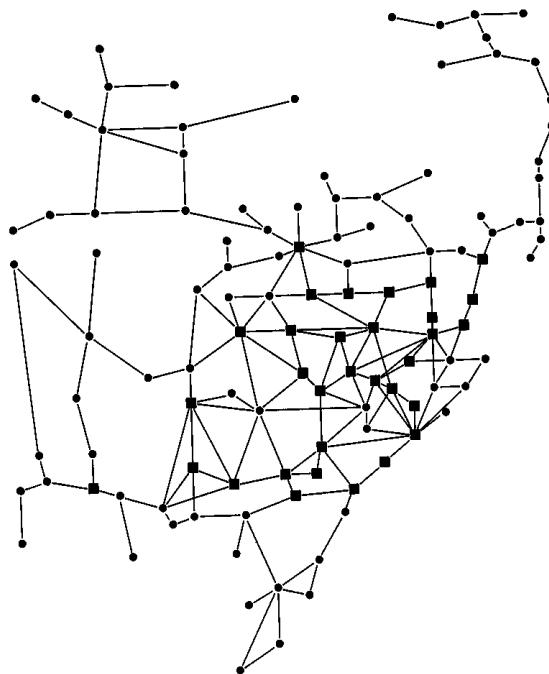


Fig. 4. Original graph of LATADL-problem.

LATA1 is a 2ECON problem, while the other six instances are 2NCON problems. All optimum solutions of the 2ECON versions turned out to satisfy all node-survivability constraints and thus were optimum solutions of the original 2NCON problems — with one exception. In LATA5L one node is especially attractive because many edges with low cost lead to it. This node is an articulation node of the optimum 2ECON solution. In the following, LATA5LE is the 2ECON version of problem LATA5L.

Table 2 contains some data about the performance of our code on the eight test instances.

We think that it is worth noting that each of these real problems, typical in size and structure, can be solved on a 28-MIPS machine in less than thirty seconds including all input and output routines, drawing the solution graph, branch and cut, etc.

A detailed analysis of the running times of the cutting plane phase is given in Table 3. All times reported are in percent of the total running time (without the branch & cut phase). The last column TT\RED shows the running times of the cutting plane phase of our algorithm applied to the full instances on the original graphs (without reduction by preprocessing). By comparing the last two columns, one can clearly see that substantial running time reductions can be achieved by our preprocessing algorithms on the larger problems.

A structural analysis of the optimum solutions produced by our code revealed that — except for LATADSF, LATA5LE, and LATA1 — the optimum survivable networks consist of a long cycle (spanning all nodes of type 2 and some nodes

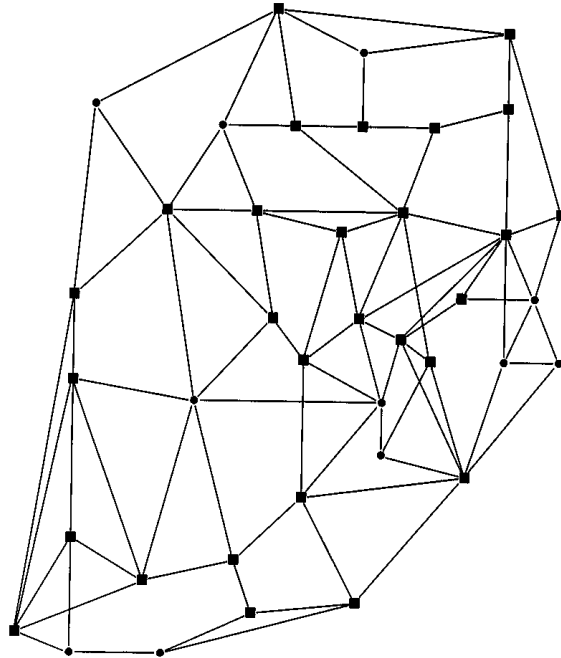


Fig. 5. Reduced graph of LATADL-problem.

Table 2

Performance of branch & cut on LATA problems

Problem	IT	P	NP	RC	C	COPT	GAP	T	BN	BD	BT
LATADMA	12	65	3	7	1489	1489	0	1			
LATA1	4	73	0	1	4296	4296	0	1			
LATASS	4	76	0	0	4739	4739	0	1			
LATASLE	7	120	0	0	4574	4574	0	1			
LATA5L	19	155	12	0	4679	4726	0.99	2	4	2	4
LATADSF	7	43	0	0	7647	7647	0	1			
LATADS	17	250	0	4	7303.60	7320	0.22	4	28	9	17
LATADL	14	182	0	28	7385.25	7400	0.20	3	32	10	21

IT = number of iterations (= calls to the LP-solver); P = number of partition inequalities (6.4) used in the cutting plane phase; NP = number of node partition inequalities (6.5) used in the cutting plane phase; RC = number of lifted r -cover inequalities (6.9) used in the cutting plane phase; C = value of the optimum solution after termination of the cutting plane phase; COPT = optimum value; GAP = $100 \times (\text{COPT} - C) / \text{COPT}$ (= percent relative error at the end of the cutting plane phase); T = total running time including input, output, preprocessing, etc., of the cutting plane phase (not including branch & cut), in rounded seconds; BN = number of branch & cut nodes generated; BD = maximum depth of the branch & cut tree; BT = total running time of the branch & cut algorithm including the cutting plane phase in seconds.

of type 1) and several branches connecting the remaining nodes of type 1 to the cycle. The optimum solution of the LATADL instance is shown in Figure 6, with the 2-connected part (the long cycle) drawn bold.

Table 3
Relative running times of cutting plane algorithm on LATA problems

Problem	PT (%)	LPT (%)	CT (%)	MT (%)	TT (s)	TT\RED (s)
LATADMA	2.0	39.2	41.2	17.6	1	1
LATA1	3.8	34.6	34.6	26.9	1	4
LATA5S	3.8	34.6	34.6	26.9	1	1
LATA5LE	0.0	42.9	41.1	16.1	1	1
LATA5L	0.7	37.1	55.2	7.0	2	5
LATADSF	2.1	21.3	57.4	19.2	1	4
LATADS	0.0	44.7	49.0	6.4	4	17
LATADL	1.0	26.3	66.2	6.5	3	18

PT = time spent in the preprocessing phase; CT = time spent in the separation routines; LPT = time used by the LP-solver; MT = miscellaneous time for input, output, drawing, etc.; TT = total time; TT\RED = total time of the algorithm when applied to the original instance without prior reduction by preprocessing.

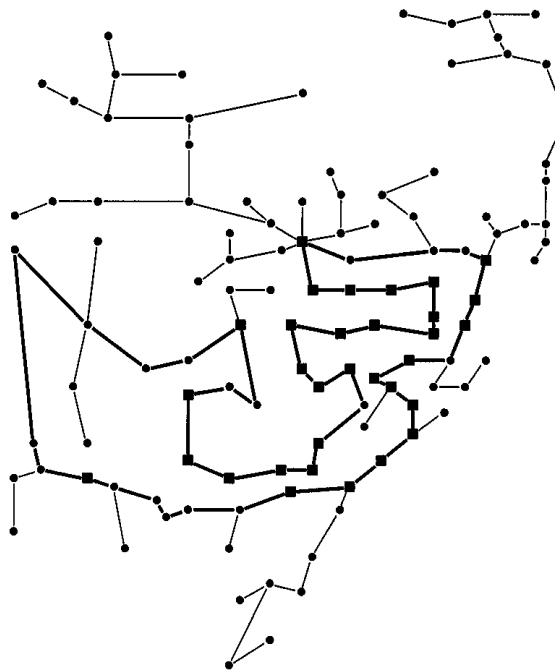


Fig. 6. Solution of LATADL-problem.

From the view of a telephone network designer, a long cycle connecting all nodes of type 2 is not a desirable feature in a communication network because routing paths are very long, resulting in delays, and because each link has to carry a high traffic load, resulting in high costs for terminal electronics, multiplexers, etc. But since the network installation costs form part of the whole network cost, the lowest network installation cost (as found by our algorithm) provides a lower bound for

Table 4
Comparison of heuristic values with optimal values

Problem	COPT	CHEUR	GAP
LATADMA	1489	1494	0.34
LATA1	4296	4296	0
LATA5S	4739	4739	0
LATA5LE	4574	4574	0
LATA5L	4726	4794	1.44
LATADSF	7647	7727	1.05
LATADS	7320	7361	0.56
LATADL	7400	7460	0.81

The whole network cost, and the subgraph minimizing the installation cost could be modified, e.g., by adding some more links of low cost, to produce a network with shorter routing paths between each pair of nodes. This is the design approach taken in the software package distributed by Bellcore [1988]; first a survivable network topology of low cost is computed by heuristics, then this topology is modified to account also for costs associated with the expected traffic in the network.

We ran a few tests on randomly generated problems of higher density and 50–100 nodes. Here our code did perform reasonably well but not as well as on sparse problems. (That is not of great importance, since our goal was to solve real-world problems and not random problems.) More serious is a dramatic increase in running time when many nodes of type 0 are added, as is the case in the ship problem treated in the next section. But the problems that we address here mainly and that come up in the design of fiber optic telephone networks have very few nodes of type 0, if any.

Another motivation for our work was to find out how well the heuristics of Monma & Shallcross [1989] described in Section 4 perform. It turned out that they do very well. Table 4 compares the values CHEUR of the solutions produced by the heuristics with the optimum values COPT computed by our code. The percent relative error GAP ($= 100 \times (\text{CHEUR} - \text{COPT})/\text{COPT}$) is always below 1.5%. In three cases the heuristics found an optimum solution. This result definitely justifies the present use of these heuristics in practice. We note that these heuristics are very fast, typically taking only a few seconds on a IBM PC/AT.

7.4. Computational results for high-connectivity problems

At present, we have a first preliminary version of a code for solving survivability problems with higher connectivity requirements. In order to test our code for general k NCON problems, we first used a set of random problems. Later, we also obtained test data for a real-world 3NCON problem, which arose in the design of a communication network on a ship. Both types of test problems have their 'drawbacks', however. The random problems turned out to be too easy (most of them were already solved in the first iteration), and the ship problem confronted

us with so many new difficulties (with respect to space, running time, and quality of solutions) that we have to redesign our separation strategies completely to solve variants of the ship problem to optimality.

7.4.1. Random problems

We first report about our computational results on random k ECON problems. We used the same set of random data as Ko & Monma [1989] used for their high-connectivity heuristics. So we will be able to compare results later.

The test set of Ko & Monma 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 next table reports the number of iterations (minimum and maximum) and the average time 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 (s)		
	$k = 3$	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

All problems except one 3ECON instance on 20 nodes were solved in the cutting plane phase. In fact 20 of the 30 problems were already solved in the first iteration with the initial LP (16). For the instances not solved in the first iteration, at most four lifted r -cover inequalities (15) had to be added to obtain the optimal solution. Except for one 3ECON instance, no partition inequalities were added. So, the average solution time is mainly the solution time for the first LP.

All optimal solutions for the k ECON problems were at the same time feasible and hence optimal for the corresponding k NCON problems, except the one 3ECON problem which could not be solved in the cutting plane phase. There the optimal solution (obtained by branch & cut) is 3-edge connected, but not 3-node connected.

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 & Monma did not perform quite as well as the low-connectivity heuristics, but still 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 11% error (taken over all problems).

7.4.2. Ship 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

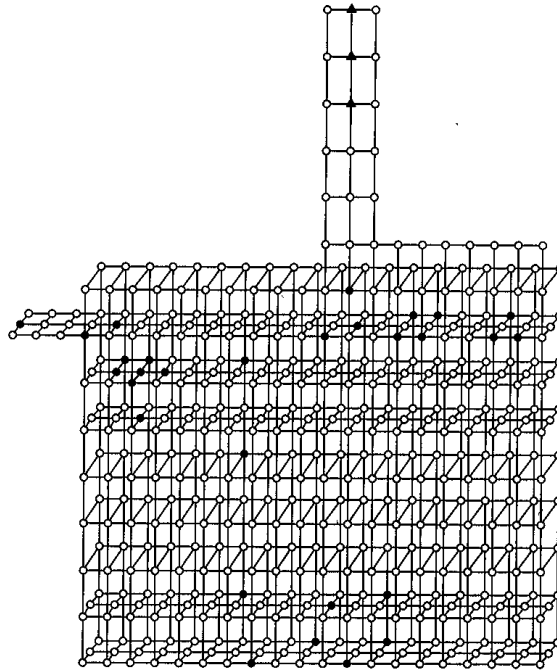


Fig. 7. Grid graph of the ship problem.

connects locations on a military ship containing various communication systems. The reason for demanding high survivability of this network is obvious.

The problem of finding a high-connected network topology minimizing the cable installation cost can be formulated as a 3NCON 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, which is depicted in Figure 7. The problem to be solved in this graph is a 3NCON problem with the following node types and costs.

Of the grid's 494 nodes, only 33 are of nonzero type, called 'special nodes'. They are drawn by filled circles or triangles. The 33 special nodes symbolize the various communication systems to be interconnected by the network. To evaluate the dependence of network topology cost on the required survivability, the ship problem appears in three different versions depending on the node types of the 33 special nodes. The three nodes depicted by triangles in the tower of the ship always have type 3, the other 29 special nodes are all given either type 1, type 2, or type 3. We call the three resulting versions of the ship problem 'ship13', 'ship23', and 'ship33', respectively. The remaining 461 nodes are nodes of type 0. They represent possible fiber junction boxes where the fiber cable may be routed.

The cost structure is highly regular. The costs are proportional to the distances between nodes, with the feature that horizontal distances are much higher than vertical distances. (The grid shown in Figure 7 has been scaled. Also, contrary to

the graphical representation, the horizontal layers do not always have the same distance from each other.) With this cost structure, it is much cheaper to route vertically than horizontally. Since there exist many shortest paths between any two nodes, there will also exist many optimum solutions to the survivable network problem. So the problem is highly degenerate. Degeneracy together with the size of the ship problem caused us to run into difficulties. In fact, when we first applied our code to the 'ship13' problem, with the initial LP consisting only of the degree constraints for the special nodes, the fractional solutions did not get connected for a long time.

Our first idea was to heuristically reduce the size of the problem in some way. Unfortunately, none of the decomposition techniques described earlier applied, except at the tower of the ship, where nodes of type 3 are separated by a cut of size 3. We cut out some of the 'unnecessary' nodes of type 0 in the lower left and right hand corner of the grid, and also deleted some of the horizontal layers of the grid containing only nodes of type 0.

It is not obvious at all that corners of a grid may be cut out and layers may be deleted without affecting the optimum objective function value of the problem. We could prove such a result only for Steiner tree problems (1NCON problems), not for 2NCON or 3NCON problems. But nevertheless, we used these reductions heuristically to cut down problem sizes in the hope that some optimal solution of the original graph is still contained in the reduced graph. For the 'ship23' problem, the optimal solution of the reduced problem turned out to be optimal for the non-reduced problem, too.

Figure 8 shows the reduced graph of the 'ship13' problem. The result of the reductions can be seen from Table 5, whose columns list, from left to right, the problem names, and, for the original ship graph and the reduced ship graphs, the number of nodes of type 0, 1, 2, and 3, the total number of nodes and the total number of edges/number of forced edges. The forced edges are those edges contained in some cut of size 3 separating two nodes of type 3, which must be contained in any feasible solution.

An optimal solution for the reduced 'ship23' problem is shown in Figure 9.

Table 5 shows that the reductions are enormous, yet there are still many more nodes of type 0 than nodes of nonzero type in each problem.

When we applied our code to the reduced graphs, the fractional solutions still looked frequently like paths beginning at some special node and ending in some

Table 5
Sizes of ship problems

Problem	Original graph					Reduced graph				
	0	1	2	3	Edges	0	1	2	3	Edges
ship13	461	30	0	3	1096/0	128	28	0	3	325/3
ship23	461	0	30	3	1096/0	249	0	30	3	607/3
ship33	461	0	0	33	1096/0	300	0	0	33	719/9

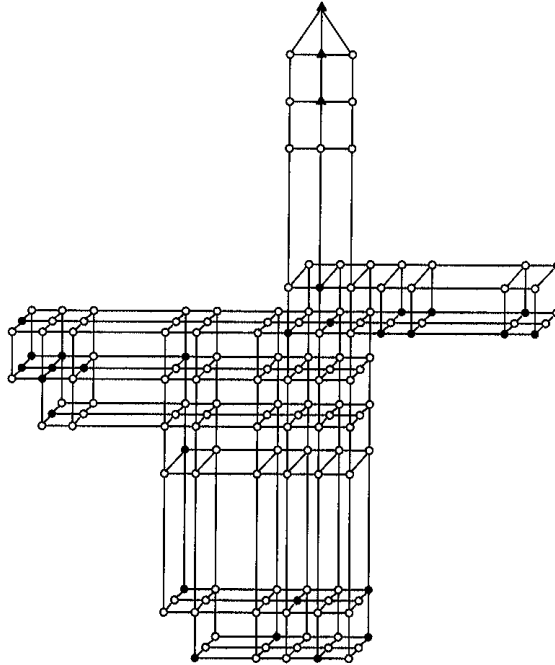


Fig. 8. Reduced grid graph of the 'ship13' problem.

node of type 0. To cure this problem, we made use of the following type of inequalities.

$$x(\delta(v) \setminus \{e\}) \geq x_e$$

for all nodes v of type 0 and all $e \in \delta(v)$. These inequalities (we call them *con0 inequalities*) describe algebraically that nodes of type 0 do not have degree 1 in an edge-minimal solution. This is not true for all survivable networks, but it is true for the optimum solution if all costs are positive. So, although these inequalities are not valid for the k NCON polytope, we used them to force the fractional solutions into the creation of longer paths. Another trick to obtain better starting solutions was to use cuts of a certain structure in the initial LP.

Table 6 gives some preliminary computational results of our cutting plane algorithm on the three reduced and not reduced versions of the ship problem.

Although Table 6 shows that the code is still rather slow, it could at least solve two of the ship problems. In order to obtain better results and running times, some more research must be done, especially on finding better starting solutions, devising faster separations heuristics that exploit the problem structure, and, maybe, inventing new classes of inequalities for high-connectivity problems. The table also shows that the speedup on the (heuristically) reduced problems was significant.

Table 7 shows the percentage of time spent in the different routines.

We do not understand yet why our code solves the ship23 problem rather easily and why there is still a gap after substantial running time of our cutting plane algorithms for the ship33 problem. Probably, the 'small' changes of a

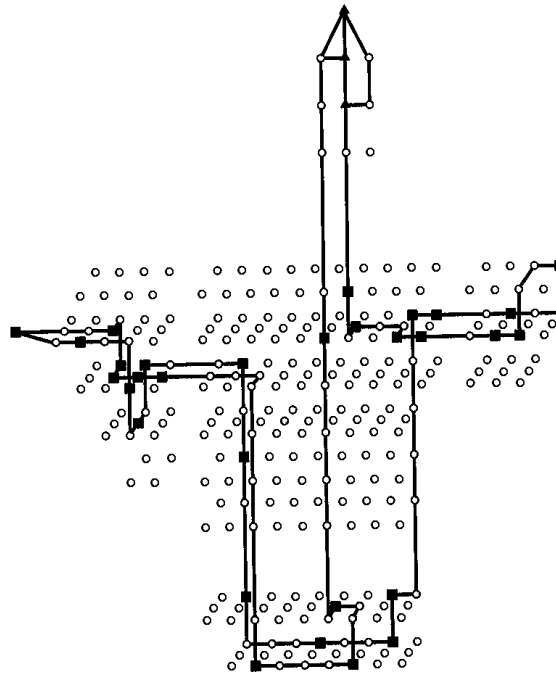


Fig. 9. Optimum solution of reduced 'ship23' problem.

Table 6

Performance of cutting plane algorithm on ship problems

Problem	VAR	IT	PART	RCOV	LB	UB	GAP (%)	Time (min:s)
ship13	1088	3252	777261	0	211957.1	217428	2.58	10122:35
ship23	1088	15	4090	0	286274	286274	0	27:20
ship33	1082	42	10718	1	461590.6	483052	4.64	55:26
ship13red	322	775	200570	0	217428	217428	0	426:47
ship23red	604	12	2372	0	286274	286274	0	1:54
ship33red	710	40	9817	0	462099.3	483052	4.53	34:52

Problem = problem name, where 'red' means reduced; VAR = number of edges minus number of forced edges; IT = number of LPs solved; PART = number of partition inequalities added; RCOV = number of r -cover inequalities added; LB = lower bound (= optimal LP value); UB = upper bound (= heuristic value); GAP = $(UB - LB)/LB$.

few survivability requirements result in more dramatic structural changes of the polyhedra and thus of the inequalities that should be used. It is conceivable that our code has to be tuned according to different survivability requirements settings. We should mention that we did not attempt to solve ship13 and ship33 by entering the branching phase of our code. The gaps are not small enough yet for the enumerative stage to have a decent perspective. Further details of our attempts to solve network design problems with higher connectivity requirements can be found in Grötschel, Monma & Stoer [1992c].

Table 7
Relative running times on ship problems

Problem	PT (%)	LPT (%)	CT (%)	MT (%)	Time (min:s)
ship13	0.0	75.6	23.9	0.5	10122:35
ship23	0.0	13.1	86.4	0.4	27:20
ship33	0.0	31.2	68.2	0.6	55:26
ship13red	0.0	68.5	30.1	1.4	426:47
ship23red	0.1	39.2	58.6	1.9	1:54
ship33red	0.0	41.1	58.4	0.5	34:52

Problem = problem name where 'red' means reduced;
PT = time spent for reduction of problem; LPT = time
spent for LP solving; CT = time spent for separation;
MT = time on miscellaneous items, input, output, etc.

Summarizing our computational results, we can say that for survivability problems with 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. But for applications in the area of telephone network design, where problem instances typically are of moderate size and contain not too many nodes of type 0, our approach produces very good lower bounds and even optimum solutions in a few minutes. This work is a good basis for the design of a production code for the 2ECON and 2NCON problems coming up in fiber optic network design and a start towards problems with higher and more varying survivability requirements and larger underlying graphs.

8. Directed variants of the general model

There are many possible variants of the general model described in Section 3 for the design of networks with connectivity constraints. A natural variant is to consider networks with directed links. As we will see below, there are practical and theoretical reasons for considering survivability in directed graphs.

8.1. Survivability models for directed networks

In order to model directed links, we let $D = (N, A)$ denote a *directed graph* consisting of a set V of *nodes* (just as in the undirected case) and a set A of *directed arcs*. Each arc $a = (u, v) \in A$ represents a link directed from node u to node v . For example, this could model certain communications facilities that allow only the one-way transfer of information. Of course, there may be arcs directed each way between any given pair of nodes. Each arc $a \in A$ has a nonnegative *fixed cost* c_a of establishing the link connection. The directed graph may have parallel arcs (in each direction). As before, the cost of establishing a network consisting of a subset $B \subseteq A$ of arcs is the sum of costs of the individual links contained in B .

The goal is to build a minimum-cost network so that the required survivability conditions are satisfied.

The survivability requirements demand that the network satisfy the same types of edge and node connectivity requirements as in the undirected case. We simply replace the notion of an undirected path by a directed one. The previous definitions and model formulations are essentially unchanged.

The problem of designing a survivable directed network has not received as much attention in the literature as the undirected case. We briefly summarize some recent efforts along these lines.

Dahl [1991] has given various formulations for the directed survivable network design problem with arc connectivity requirements. He mainly studies the bi-Steiner problem, which is the problem of finding a minimum-cost directed subgraph that contains two arc-disjoint paths from a given root to each node of a set of terminal nodes. This problem has applications in the design of hierarchical subscriber networks, see Lorentzen & Moseby [1989].

Chopra [1992] modeled a directed version of the 2ECON problem, which becomes the 'undirected' 2ECON problem after 'projection' into a lower-dimensional space. He showed that all partition inequalities and further inequalities can be generated by the projection of certain directed cut inequalities. Chopra's model can be generalized to higher edge connectivity requirements, as shown below.

8.2. Projection

The last remarks show that directed versions of the k ECON and k NCON problems are not only interesting in their own right, but they are sometimes also useful in solving their undirected counterparts. We will illustrate this now by pointing out the value of projections. For many combinatorial problems, good polyhedral descriptions can be obtained by transferring the original problem into higher dimensions, that is, by formulating it with additional (auxiliary) variables, which may later be projected away. This was done successfully for the 2-terminal Steiner tree problem in directed graphs, see Ball, Liu & Pulleyblank [1987]. There the formulation with auxiliary variables contains a polynomial number of simple constraints, which by projection are turned into an exponential number of 'weird' constraints. The general idea of projection was described by Balas & Pulleyblank [1983].

For the 2ECON problem, Chopra [1992] has found a formulation in directed graphs using $2|E|$ integer variables and directed cut constraints, which he called the DECON problem, see (17) below. The directed cut constraints, see (17i), used in the formulation of the DECON problem have the advantage that they can be separated in polynomial time, whereas the separation of the inequalities appearing in our undirected 2ECON problem is NP-hard.

Projection of the directed cut constraints and nonnegativity constraints of the DECON problem gives a new class of inequalities for the 2ECON problem (we call these Prodon inequalities) which contain as a subclass the partition

inequalities (10). For the Steiner tree problem, where $r_v \in \{0, 1\}$, these new inequalities have been found by Prodon [1985]. In the following we show how the Prodon inequalities are derived from the DECON model by projection.

In order to do this, we must first introduce some terminology. Let a graph $G = (V, E)$ and node types $r_v \in \{0, 1, 2\}$ be given, where at least two nodes are of highest (positive) node type. This may either be a 2ECON or a 1ECON problem. From G we construct a directed graph $D = (V, A)$ by replacing each undirected edge ij with two directed edges (i, j) and (j, i) . Furthermore, we pick some node $w \in V$ of highest node type. Let $\delta^-(W)$ be the set of arcs directed into node set W . If (x, y) is a solution to the following system of inequalities (where $x \in \mathbf{Z}^E$ and $y \in \mathbf{Z}^A$),

$$\begin{aligned}
 \text{(i)} \quad & y(\delta^-(W)) \geq 1 \text{ for all } W \subseteq V, \emptyset \neq W \neq V, \text{ with} \\
 & \quad \text{con}(W) = 2 \text{ (or } r(W) = 1 \text{ and } w \notin W); \\
 \text{(ii)} \quad & y_{(i,j)} \geq 0 \text{ for all } (i, j) \in A; \\
 \text{(iii)} \quad & y_{(i,j)} \text{ integral for all } (i, j) \in A; \\
 \text{(iv)} \quad & -y_{(i,j)} - y_{(j,i)} + x_{ij} = 0 \text{ for all } ij \in E; \\
 \text{(v)} \quad & x_{ij} \leq 1 \text{ for all } ij \in E.
 \end{aligned} \tag{17}$$

then the integer vector x is feasible for the 2ECON problem, and vice versa: if some integer vector x is feasible for the 2ECON problem, then an integer vector y can be found so that (x, y) satisfies (17i)–(17v). So the projection of system (17) onto x -variables gives a formulation of the 2ECON problem. (Originally, Chopra considered this system without the upper bound constraints.) If no node is of type 2, a feasible vector y is just the incidence vector of a subgraph of D containing a Steiner tree rooted at w . If all nodes are of type 2, then y is the incidence vector of a strongly connected directed subgraph of D ('strongly connected' means that between each distinct pair s, t of nodes there exists a directed (s, t) -path and a directed (t, s) -path).

Without the integrality constraints (17iii) and upper bound constraints (17v), we obtain a relaxation, which, after projection onto x -variables, gives a relaxation of the 2ECON problem. The projection works as follows. Let us define

(1) \mathcal{F} as the set of those $W \subseteq V$ that appear in the formulation of inequalities (17i),

(2) $b_W \geq 0$ as the variables assigned to each inequality (17i) for $W \in \mathcal{F}$,

(3) $a_{ij} \in \mathbf{R}$ as the variables assigned to each equation (17iv) for $ij \in E$,

(4) $s(\mathcal{F}; b; i; j)$ as the sum of b_W over all $W \in \mathcal{F}$ with $i \in W$ and $j \notin W$, and

(5) C as the cone of variables $a \in \mathbf{R}^E$ and $b := (b_W)_{W \in \mathcal{F}}$ satisfying

$$\begin{aligned}
 a_{ij} &\geq s(\mathcal{F}; b; i; j), \text{ for all } ij \in E, \\
 a_{ij} &\geq s(\mathcal{F}; b; j; i), \text{ for all } ij \in E, \\
 b &\geq 0.
 \end{aligned}$$

If $(a, b) \in C$, and if all inequalities of type (17i) and all inequalities of type (17iv) are added with coefficients b_W and a_{ij} respectively, then we obtain an

inequality

$$\sum_{(i,j) \in A} u_{(i,j)} y_{(i,j)} + \sum_{ij \in E} a_{ij} x_{ij} \geq \sum_{W \in \mathcal{F}} b_W,$$

where the $u_{(i,j)}$ are non-positive coefficients of the variables $y_{(i,j)}$. In fact, C was defined exactly in such a way, that the $u_{(i,j)}$ are non-positive. The above inequality is valid for the system given by all inequalities (17i), (17ii), and (17iv). Since $y \geq 0$,

$$\sum_{ij \in E} a_{ij} x_{ij} \geq \sum_{W \in \mathcal{F}} b_W,$$

is valid for $2ECON(G; r)$. It can also be proved with the general projection technique of Balas & Pulleyblank [1983] that

$$\sum_{ij \in E} a_{ij} x_{ij} \geq \sum_{W \in \mathcal{F}} b_W \text{ for all } (a, b) \in C, \tag{18}$$

$$x \geq 0$$

is exactly the projection of system (17i), (17ii) and (17iv) onto the x -variables. Not all $(a, b) \in C$ are needed in the formulation of (18). The following system is clearly sufficient to describe the projection of (17i), (17ii) and (17iv) onto x -variables:

$$\begin{aligned} \text{(i)} \quad & \sum_{ij \in E} a_{ij} x_{ij} \geq \sum_{W \in \mathcal{F}} b_W \text{ for all } b \geq 0 \text{ and} \\ & a_{ij} := \max\{s(\mathcal{F}; b; i; j), s(\mathcal{F}; b; j; i)\}, \tag{19} \\ \text{(ii)} \quad & x \geq 0. \end{aligned}$$

We call inequalities (19i) *Prodon inequalities* (induced by b), because this class of inequalities was discovered by Prodon [1985] for $1ECON(G; r)$.

The class of Prodon inequalities properly contains the class of partition inequalities (10). Namely, a partition inequality

$$x[W_1 : \dots : W_p] \geq \begin{cases} p & \text{if at least two } W_i \text{ contain nodes of type 2} \\ p - 1 & \text{otherwise} \end{cases}$$

(where W_1, \dots, W_p is a partition of V into p node sets with $r(W_i) \geq 1$) can also be written as a Prodon inequality, if b_W is set to 1 for all W_i that are in \mathcal{F} and $b_W := 0$ for all other sets in \mathcal{F} . By definition of \mathcal{F} , if at least two sets W_i contain nodes of type 2, then $W_i \in \mathcal{F}$ for all W_i , and if only one set, say W_p , contains nodes of type 2 (and therefore the 'root' w), then W_1, \dots, W_{p-1} are in \mathcal{F} , but W_p is not. This explains the differing right-hand sides in both cases.

But not every facet-defining Prodon inequality is also a partition inequality. For instance, the inequality depicted in Figure 10 is not a partition inequality, but can be written as a Prodon inequality induced by $b_W := 1$ for the sets $\{1\}, \{2\}, \{5\}, \{7\}, \{3, 5, 6\}, \{4, 6, 7\}$, and $b_W := 0$ for all other sets W in \mathcal{F} . So the coefficients on all depicted edges are 1, and the right-hand side is 6. Here, nodes 1 and 2 are nodes of type 2; nodes 5 and 7 are nodes of type 1; all others are of type 0. The Prodon inequality of Figure 10 can be proved to be facet-defining for $2NCON(G; r)$, where G consists exactly of the depicted nodes and edges.

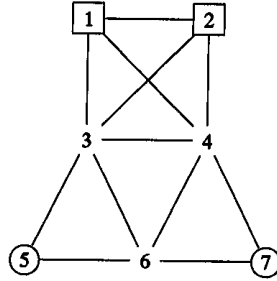


Fig. 10. Prodon inequality.

We show in the following remark that no Prodon inequality except the cut inequalities are facet-defining if there are no nodes of type 1.

Remark 2. *If (G, r) is an instance of the 2ECON problem, where node types r_v only take values 0 and 2 for all $v \in V$, then no Prodon inequalities except the cut constraints define facets of $2ECON(G; r)$.*

Proof. Let $\sum_{ij} a_{ij} x_{ij} \geq \sum_{W \in \mathcal{F}} b_W$ be a Prodon inequality. By definition,

$$a_{ij} \geq \frac{1}{2} s(\mathcal{F}; b; i; j) + \frac{1}{2} s(\mathcal{F}; b; j; i),$$

which is the same as $1/2$ times the sum of all b_W over $W \in \mathcal{F}$ with $ij \in \delta(W)$. Therefore,

$$a^T x \geq \frac{1}{2} \sum_{W \in \mathcal{F}} b_W x(\delta(W)).$$

Since $x(\delta(W)) \geq \text{con}(W) = 2$ for all $W \in \mathcal{F}$, this expression is at least $\sum_{W \in \mathcal{F}} b_W$ for all $x \in 2ECON(G; r)$. So our Prodon inequality is implied by the sum of some cut inequalities, and must itself be a cut inequality, if it is to be facet-defining. \square

The projection technique applied to the Steiner tree polytope is not new. Goemans & Myung [1993] list various formulations of the Steiner tree problem, all of which use auxiliary variables, among them system (17) [without (v)]. They show that upon projection to variables $x \in \mathbf{R}^E$ all these models generate the same inequalities. Goemans [1994b] investigates, again for the Steiner tree polytope, facet properties of a subclass of inequalities obtained by such a projection, which are, in fact, a subclass of the class of Prodon inequalities.

8.2.1. Higher connectivity requirements

The DECON model can be generalized to higher connectivity requirements, when node types are in $\{0, 1, 2, 4, 6, \dots\}$. The directed model in this case requires $(1/2) \min\{r_u, r_v\}$ directed arc-disjoint (u, v) -paths between each pair of nodes u, v whose node types are at least 2, and one directed path from a specified root of highest node type to each node of type 1. This does appropriately model

the undirected k ECON problem, because of a theorem of Nash-Williams [1960], which says that undirected graphs containing r_{uv} edge-disjoint paths between each pair of nodes u and v can be oriented in such a way, that the resulting directed graph contains, between each u and v , $\lfloor r_{uv}/2 \rfloor$ arc-disjoint paths. The inequalities resulting from the projection of directed cut inequalities in this model do *not* generalize all partition inequalities of type (10) when $k \geq 4$.

8.2.2. Separation of Prodon inequalities

We close this section by observing that the separation problem for Prodon inequalities can be performed in polynomial time. The separation algorithm also makes use of projection and works in the same way as one iteration of Benders' decomposition method, see Benders [1962]. This observation is meant to show that projection is not only of theoretical value but also of computational interest.

Suppose a point x^* with $0 \leq x^* \leq 1$ is given, for which it has to be decided whether there is a Prodon inequality violated by this point or not. This can be decided by solving the following LP derived from (17):

$$\begin{aligned} & \min z \\ & \text{subject to} \\ & \text{(i) } y(\delta^-(W)) + z \geq 1 \quad \text{for all } W \in \mathcal{F}; \\ & \text{(ii) } y_{(i,j)} \geq 0 \quad \text{for all } (i,j) \in A; \\ & \text{(iii) } -y_{(i,j)} - y_{(j,i)} - zx_{ij}^* = -x_{ij}^* \quad \text{for all } ij \in E; \\ & \text{(iv) } z \geq 0. \end{aligned} \tag{20}$$

This LP has the feasible solution $y = 0$ and $z = 1$. If its optimal value is 0, and y^* is an optimal solution, then (x^*, y^*) is feasible for the system (17), hence x^* satisfies all Prodon inequalities (by the projection result). If the optimal value is non-zero, then the optimal dual variables b_W for inequalities (20i) and a_{ij} for equations (20iii) define a Prodon inequality violated by x^* . More explicitly, the optimal dual variables b_W ($W \in \mathcal{F}$) and $a \in \mathbf{R}^E$ satisfy

$$\begin{aligned} -a_{ij} + \sum_{W \in \mathcal{F}: i \in W, j \notin W} b_W &\leq 0 \quad \text{for all } (i,j) \in A, \\ -a_{ij} + \sum_{W \in \mathcal{F}: j \in W, i \notin W} b_W &\leq 0 \quad \text{for all } (i,j) \in A, \\ -a^T x^* + \sum_{W \in \mathcal{F}} b_W &> 0. \end{aligned} \tag{21}$$

The first two inequalities imply that a_{ij} is at least the maximum of $s(\mathcal{F}; b; i; j)$ and $s(\mathcal{F}; b; j; i)$ for each $ij \in E$. This implies that a and b induce the Prodon inequality

$$\sum_{ij \in E} a_{ij} x_{ij} \geq \sum_{W \in \mathcal{F}} b_W.$$

From the last inequality in (21) it follows that x^* violates this Prodon inequality.

The LP (20) can be solved in polynomial time, since there exist polynomial separation algorithms for the directed cut inequalities (20i). Therefore, the Prodon inequalities can also be separated in polynomial time. We have, however, not yet made use of these inequalities.

References

- Agrawal, A., Ph. Klein and R. Ravi (1991). When trees collide: An approximation algorithm for the generalized Steiner tree problem on networks. *Proc. 23rd Annu. Symp. on Theory of Computing*, pp. 134–144, May 1991.
- Baiou, M., and A.R. Mahjoub (1993). The 2-edge connected Steiner subgraph polytope of a series-parallel graph, Département d'Informatique, Université de Bretagne Occidentale, France, October 1993.
- Balas, E., and W.R. Pulleyblank (1983). The perfectly matchable subgraph polytope of a bipartite graph. *Networks* 13, 495–516.
- Ball, M.O., W.G. Liu and W.R. Pulleyblank (1987). Two-terminal Steiner tree polyhedra, Technical Report 87466-OR, University of Bonn.
- Bellcore (1988). FIBER OPTIONS: Software for designing survivable optical fiber networks, Software Package, Bell Communications Research.
- Benders, J.F. (1962). Partitioning procedures for solving mixed-variable programming problems. *Numer. Math.* 4, 238–252.
- Bienstock, D., E.F. Brickell and C.L. Monma (1990). On the structure of minimum-weight k -connected spanning networks, *SIAM J. Discrete Math.* 3, 320–329.
- Bixby, R.E. (1992). Implementing the simplex method: The initial basis, *ORSA J. Comput.* 4, 267–284.
- Bland, R.G., D. Goldfarb and M.J. Todd (1981). The ellipsoid method: a survey. *Oper. Res.* 29, 1039–1091.
- Boyd, S.C., and T. Hao (1993). An integer polytope related to the design of survivable communication networks, *SIAM J. Discrete Math.* 6(4), 612–630.
- Cai, G.-R., and Y.-G. Sun (1989). The minimum augmentation of any graph to a k -edge connected graph. *Networks* 19, 151–172.
- Cardwell, R.H., C.L. Monma and T.H. Wu (1989). Computer-aided design procedures for survivable fiber optic networks, *IEEE J. Selected Areas Commun.* 7, 1188–1197.
- Cardwell, R.H., T.H. Wu and W.E. Woodall (1988). Decreasing survivable network cost using optical switches, in: *Proc. GLOBECOM '88*, pp. 93–97.
- Chopra, S. (1992). Polyhedra of the equivalent subgraph problem and some edge connectivity problems, *SIAM J. Discrete Math.* 5(3), 321–337.
- Chopra, S. (1994). The k -edge connected spanning subgraph polyhedron, *SIAM J. Discrete Math.* 7(2), 245–259.
- Chou, W., and H. Frank (1970). Survivable communication networks and the terminal capacity matrix, *IEEE Trans. Circuit Theor.* CT-17(2), 192–197.
- Chvátal, V. (1973). Edmonds polytopes and a hierarchy of combinatorial problems, *Discrete Math.* 4, 305–337.
- Cornuéjols, G., J. Fonlupt and D. Naddef (1985). The traveling salesman problem on a graph and some related integer polyhedra. *Math. Program.* 33, 1–27.
- Dahl, G. (1991). Contributions to the design of survivable directed networks, Ph.D. Thesis, University of Oslo. Technical Report TF R 48/91, Norwegian Telecom, Research Dept., Kjeller, Norway.
- Edmonds, J. (1965). Maximum matching and a polyhedron with 0,1-vertices, *J. Res. Nat. Bur. Stand. Ser. B* 69, 125–130.
- Eswaran, K.P., and R.E. Tarjan (1976). Augmentation problems. *SIAM J. Comput.* 5(4), 653–665.

- Frank, A. (1992a). Augmenting graphs to meet edge-connectivity requirements. *SIAM J. Discrete Math.* 5(1), 25–53.
- Frank, A. (1992b) On a theorem of Mader. *Ann. Discrete Math.* 101, 49–57.
- Frank, A. (1995). *Connectivity and network flows*, in: R. Graham, M. Grötschel and L. Lovász (eds.), *Handbook of Combinatorics*, North Holland, Amsterdam, Chapter 2, to appear.
- Frank, A., and T. Jordán (1993). Minimal Edge-Coverings of Pairs of Sets, Research Institute for Discrete Mathematics, University of Bonn, Germany, June 1993.
- Frank, H., and W. Chou (1970). Connectivity considerations in the design of survivable networks. *IEEE Trans. Circuit Theor.* CT-17(4), 486–490.
- Frederickson, G.N., and J. Jájá (1982). On the relationship between the biconnectivity augmentation and traveling salesman problem. *Theor. Comput. Sci.* 19, 189–201.
- Garey, M.R., and D.S. Johnson (1979). *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco, Calif.
- Goemans, M.X. (1994a). Arborescence polytopes for series-parallel graphs, *Discrete Appl. Math.* 51(3), 277–289.
- Goemans, M.X. (1994b). The Steiner tree polytope and related polyhedra, *Math. Program.* A63(2), 157–182.
- Goemans, M.X., and D.J. Bertsimas (1993). Survivable networks, linear programming relaxations and the parsimonious property. *Math. Program.* 60(2), 145–166.
- Goemans, M.X., M. Mihail, V. Vazirani, D. Williamson (1992). An approximation algorithm for general graph connectivity problems, preliminary version. *Proc. 25th ACM Symp. on the Theory of Computing*, San Diego, CA, 1993, pp. 708–717.
- Goemans, M.X., and Y.-S. Myung (1993). A catalog of Steiner tree formulations, *Networks* 23(1), 19–28.
- Gomory, R.E., and T.C. Hu (1961). Multi-terminal network flows. *J. Soc. Ind. Appl. Math.* 9, 551–570.
- Grötschel, M., L. Lovász and A. Schrijver (1988). *Geometric algorithms and combinatorial optimization*, Springer, Berlin.
- Grötschel, M., and C.L. Monma (1990). Integer polyhedra associated with certain network design problems with connectivity constraints. *SIAM J. Discrete Math.* 3, 502–523.
- Grötschel, M., C.L. Monma and M. Stoer (1992a). Facets for polyhedra arising in the design of communication networks with low-connectivity constraints, *SIAM J. Optimization* 2, 474–504.
- Grötschel, M., C.L. Monma and M. Stoer (1992b). Computational results with a cutting plane algorithm for designing communication networks with low-connectivity constraints, *Oper. Res.* 40, 309–330.
- Grötschel, M., C.L. Monma and M. Stoer (1992c). Polyhedral and computational investigations for designing communication networks with high survivability requirements, ZIB-Preprint SC 92-24, Konrad-Zuse-Zentrum für Informationstechnik Berlin.
- Grötschel, M. and M.W. Padberg (1985). Polyhedral theory, in: E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan and D. Shmoys (eds.), *The traveling salesman problem*, Wiley, Chichester, pp. 251–305.
- Hao, J., and J.B. Orlin (1992). A faster algorithm for finding the minimum cut in a graph. *Proc. 3rd Annu ACM-SIAM-Symp on Discrete Algorithms*, Orlando, Florida, pp. 165–174.
- Harary, F. (1962). The maximum connectivity of a graph. *Proc Nat. Acad. Sci., USA* 48, 1142–1146.
- Hsu, T.S. (1992). On four-connecting a triconnected graph (extended abstract). *Proc. 33rd Annu. IEEE Symp. on the Foundations of Computer Science*, pp. 70–79.
- Hsu, T.S., and V. Ramachandran (1991). A linear-time algorithm for triconnectivity augmentation (extended abstract). *Proc. 32nd Annu. Symp. on the Foundations of Computer Science*, pp. 548–559.
- Khuller, S., and U. Vishkin (1994). Biconnectivity approximations and graph carvings. *J. ACM* 41(2), 214–235.
- Ko, C.-W., and C.L. Monma (1989). Heuristic methods for designing highly survivable communication networks, Technical report, Bell Communications Research.

- Kolar, D.J., and T.H. Wu (1988). A study of survivability versus cost for several fiber network architectures. *Proc. ICC '88*, pp. 61–66.
- Lawler, E. (1976) *Combinatorial Optimization: Networks and Matroids*. Holt, Reinhart and Winston, New York.
- Lawler, E.L., J.K. Lenstra, A.H.G. Rinnooy Kan and D. Shmoys (1985). *The traveling salesman problem*, Wiley, Chichester.
- Ling, F., and T. Kameda (1987). Complexity of graph connectivity functions, Technical Report, School of Computing Science, Simon Fraser University, Burnaby, British Columbia.
- Lorentzen, R., and H. Moseby (1989). Mathematical models and algorithms used in the subscriber network planning tool ABONETT, Norwegian Telecommunications Research Dept., TF-report 66/89.
- Lovász, L. (1976). On some connectivity properties of Eulerian graphs. *Acta Math. Acad. Sci. Hung.* 28, 129–138.
- Lovász, L., and M.D. Plummer (1986). *Matching theory*. Annals of Discrete Mathematics, Vol. 29, North-Holland, Amsterdam.
- Mader, W. (1978). A reduction method for edge-connectivity in graphs. *Ann. Discrete Math.* 3, 145–164.
- Mahjoub, A.R. (1994). Two edge connected spanning subgraphs and polyhedra, *Math. Program.* A64(2), 199–208.
- Margot, F., A. Prodon and Th.M. Liebling (1994). Tree polyhedron on 2-trees, *Math. Program.* A63(2), 183–191.
- Monma, C.L., B.S. Munson and W.R. Pulleyblank (1990). Minimum-weight two-connected spanning networks. *Math. Program.* 46, 153–171.
- Monma, C.L., and D.F. Shallcross (1989). Methods for designing communication networks with certain two-connected survivability constraints. *Oper. Res.* 37, 531–541.
- Nagamochi, H., and T. Ibaraki (1992). Computing edge-connectivities in multigraphs and capacitated graphs. *SIAM J. Discrete Math.* 5(1), 54–66.
- Naor, D., D. Gusfield and C. Martel (1990). A fast algorithm for optimally increasing the edge-connectivity. *Proc. 31st Annu. Symp. on the Foundation of Computer Science*, pp. 698–707.
- Nash-Williams, C.St.J.A. (1960). On orientations, connectivity, and odd vertex pairings in finite graphs. *Can. J. Math.* 12, 555–567.
- Newark Star Ledger (1987). Damage to fiber cable hinders phone service, September 22, 1987.
- Newark Star Ledger (1988a). Cable snaps, snags area phone calls, February 26, 1988.
- Newark Star Ledger (1988b). Phone snafu isolates New Jersey; long-distance cable snaps, November 19, 1988.
- New York Times (1988). Phone system feared vulnerable to wider disruptions of service, May 26, 1988.
- New York Times (1989). Experts say phone system is vulnerable to terrorists, February 8, 1989.
- Padberg, M.W., and M.R. Rao (1982). Odd minimum cut sets and b -matchings. *Math. Oper. Res.* 7, 67–80.
- Papadimitriou, C.H., and K. Steiglitz (1982). *Combinatorial optimization: algorithms and complexity*, Prentice-Hall, Englewood Cliffs, N.J.
- Papadimitriou, C.H., and M. Yannakakis (1982). The complexity of facets and some facets of complexity *J. Assoc. Comput. Mach.* 29, 285–309.
- Prodon, A., Th.M. Liebling and H. Gröfiin (1985). Steiner's problem on two-trees, Technical Report RO-830315, École Polytechnique Fédérale de Lausanne, Switzerland.
- Prodon, A. (1985). A polyhedron for Steiner trees in series-parallel graphs, Technical Report, École Polytechnique Fédérale de Lausanne, Switzerland.
- Pulleyblank, W.R. (1989). Polyhedral combinatorics, in: G.L. Nemhauser, A.H.G. Rinnooy Kan and M.J. Todd (eds.), *Optimization*, Handbook in Operations Research and Management Science, Vol. 1, North-Holland, Amsterdam, 371–446.
- Steiglitz, K., P. Weiner and D.J. Kleitman (1969). The design of minimum cost survivable networks. *IEEE Trans. Circuit Theor.* 16, 455–460.

- Stoer, M. (1992). Design of survivable networks, Ph.D. thesis, University of Augsburg. *Lecture Notes in Mathematics*, Springer, Heidelberg, Vol. 1531.
- Takamizawa, K., T. Nishizeki and N. Saito (1982). Linear-time computability of combinatorial problems on series-parallel graphs. *J. Ass. Comput. Mach.* 29(3), 623–641.
- Ueno, S., Y. Kajitani and H. Wada (1988). Minimum augmentation of a tree to a k -edge-connected graph, *Networks* 18, 19–25.
- Wald, J.A., and C.J. Colbourn (1983). Steiner trees, partial 2-trees and minimum IFI networks. *Networks* 13, 159–167.
- Wall Street Journal (1988). Fire in fiber gateway sparks flight delays, problems at brokerages, May 11, 1988.
- Watanabe, T., and A. Nakamura (1987). Edge-connectivity augmentation problems. *Comput. System Sci.* 35(1), 96–144.
- Winter, P. (1985a). Generalized Steiner tree problem in Halin networks. *Proc. 12th Int. Symp. on Mathematical Programming*, MIT.
- Winter, P. (1985b). Generalized Steiner problem in outerplanar networks, *BIT* 25, 485–496.
- Winter, P. (1986). Generalized Steiner problem in series-parallel networks, *J. Algorithms* 7, 549–566.
- Winter, P. (1987). Steiner problem in networks: A survey. *Networks* 17(2), 129–167.
- Wu, T.H., and R.H. Cardwell (1988). Optimum routing in fiber network design: models and applications. *Proc. ICC '88*, pp. 251–257.
- Wu, T.H., D.J. Kolar and R.H. Cardwell (1988). Survivable network architectures for broadband fiber optic networks: model and performance comparison. *IEEE J. Lightwave Technol.* 6, 1698–1709.
- Zorpette, G. (1989). Keeping the phone lines open. *IEEE Spectrum*, June 1989, pp. 32–36.