Capacity and Survivability Models for Telecommunication Networks

Dimitris Alevras Martin Grötschel Roland Wessäly

Abstract

Designing low-cost networks that survive certain failure situations is one of the prime tasks in the telecommunication industry. In this paper we survey the development of models for network survivability used in practice in the last ten years. We indicate how algorithms integrating polyhedral combinatorics, linear programming, and various heuristic ideas can help solve real-world network dimensioning instances to optimality or within reasonable quality guarantees in acceptable running times.

The most general problem type we address is the following. Let a communication demand between each pair of nodes of a telecommunication network be given. We consider the problem of choosing, among a discrete set of possible capacities, which capacity to install on each of the possible edges of the network in order to satisfy all demands and to minimize the building cost of the network. In addition to determining the network topology and the edge capacities we have to provide, for each demand, a routing such that no path can carry more than a given percentage of the demand and no path in the routing exceeds a given length. We also have to make sure that for every single node or edge failure, a certain percentage of the demand is reroutable. Moreover, for all failure situations feasible routings must be computed.

The model described above has been develop in cooperation with a German mobile phone provider. We present a mixed-integer programming formulation of this model and computational results with data from practice.

1 Introduction and Survey

In this paper we describe a series of mathematical models that have been developed in the recent years to describe and solve various telecommunication network design problems. Along with the solution methodology the users of these models have become more sophisticated, demanding the integration of tasks into one model that have traditionally been solved in a hierarchical fashion. A typical sequence of such decisions consists, among other issues to be considered, of the choice of technologies to be used, the topological design of the network, the planning of the capacities of the network components, a decision about routing strategies, and the treatment of failure situations. Some companies, such as our partner `e-plus`, handle this complex suite of decisions in one integrated mixed-integer linear programming model, as we will describe later.

Before reviewing the literature on telecommunication network design we present a framework that enables us to classify the models.

The problems we consider have the following in common. The input consists of two graphs on the same node-set V, the **supply graph** $G = (V, E)$ and the **demand graph** $H = (V, D)$. The set V consists of the nodes of the logical transport network. In the application we will focus on, V is the set of MSC locations (MSC = Mobile Switching Center); in some cases BSC locations are included (BSC = Base Station Controller). The edge-set E of the supply graph G is the set of all physical links that may potentially be used (in the planning period). Different types of links (representing different technologies, e.g., microwave connections, copper or fiber optic cables, leased lines, etc.) are represented by **parallel edges**. The demand graph H (for the planning period) contains an edge whenever there is a positive demand between its two end nodes. For each edge $uv \in D$ of the demand graph, the value

- $d_{uv} \in \mathbb{Z}_+$ is the **communication demand** between nodes u and v.

While the characteristics of the supply graph are relatively stable (they change, e.g., with hardware and suppliers), demand predictions are based on statistics and marketing forecasts. They are altered frequently, and scenario analysis has to be made to take different possible evolvements of the market into account.

Given this as basic input there are several levels of possible network design problems.

Capacities

It may be that the network designer is only interested in the topological structure and has decided to determine link capacities in a later stage. It may also be that capacities are no issue since the standard technology supplies enough for the application in question. We label this situation **no capacities**. If capacity planning is necessary, capacity may be selected arbitrarily from a certain range, or only finitely many choices may be available. We label these possibilities **continuous capacities** and **discrete capacities**.

Survivability

It has become common to call a network **survivable** if it has been designed in such a way that the network is operational (in some sense to be made specific) even if certain network components fail. If the network components are very reliable and impact on the components from outside is unlikely, survivability may not be an issue. We say that **no survivability** is considered. A frequently used method to guarantee *topological survivability* is to require that the supply network to be designed contains, for each pair of nodes, a certain (node-pair dependent) number of node- and/or edge-disjoint paths between these nodes. Making reference to the graph theoretical background of this concept we say that **k-connectivity** is required. If routing is an integral part of the network design problem a reasonable strategy to keep the network "alive" in failure situations is to require, that for every pair of demand nodes, no path of the network carries more than a certain percentage of the total traffic between the two nodes. This concept runs under the name **diversification.** In case a model integrates capacity planning and routing, a natural variation of the k-connectivity concept is to require, that, for every pair of demand nodes and for every failure situation, a certain percentage of the traffic demand between the node pair can be routed. This concept is called **reservation**. We present in Section 2 two different reservation methods which differ in the way the rerouting is treated.

Path length

In case of very tight capacities, some connections, i.e., paths carrying traffic between demand nodes, turn out to be long. For various reasons (e.g., to reduce time delay or computer load) it may be advisable to restrict the length of communication paths. Thus there are models with and without path-length restriction.

Network Dimensioning Models: A Survey

In the last years several groups of authors considered various combinations of capacity and survivability models. We briefly review the four main directions.

No capacities and k-connectivity

The uncapacitated problem with connectivity requirements was one of the first network design problems investigated. It includes the Steiner-tree problem as a special case. Monma and Shallcross present a heuristic approach for the 2-connectivity case in [15]. Their heuristics produce good solutions for the "LATA networks" of Bellcore (with up to 116 nodes) in short running times. Theoretical investigations on the structure of optimal solutions can be found in Monma, Munson and Pulleyblank [14]. Grötschel, Monma and Stoer develop in [8, 9] a framework (based on branch&cut methods) to solve the LATA networks of Bellcore for low-connectivity ($k \leq 2$) instances to optimality. Furthermore, Stoer [17] reports that special high-connectivity problems ($k > 2$) can be solved to optimality for up to 500 nodes. A detailed survey of the work on this type of survivable networks can be found in [10].

Continuous capacities and reservation

Minoux [13] was the first to consider survivability in a generalized multicommodity-flow model with continuous capacities. He reports that instances with up to 40 nodes can be solved with an accuracy of about 5%. In cooperation with *France Telecom,* Lisser, Sarkissian and Vial [11] develop another model including non-discrete capacities and survivability. In [11] two survivability models are presented, both different from ours. In both models, part of the demand is routed in a separate network, called spare network, in case of a failure. The *local-survivability* model routes only the failing flow, and the *global-survivability* model routes only the affected demands in the spare network. Tests with up to 53 nodes are reported.

Discrete capacities and no survivability

Several models in the literature consider the installation of discrete capacities without addressing survivability issues. Moreover, these models restrict the possible capacities to multiples of one or two basic capacities. Bienstock and Günlük [5] solve ATM network design problems with real-life data for instances of up to 16 nodes to optimality. Their model includes flow costs, and the capacities can be chosen as combinations of two basic technologies (OC3 and

OC12 facilities). In another study with one basic technology Bienstock, Chopra, Günlük, and Tsai (see [4]) solve "New York area problems" with up to 15 nodes and "Norwegian Problems" with up to 27 nodes (supplied by M. Stoer) almost to optimality. Magnanti, Mirchandani, and Vachani [12] investigate the same problem without flow costs and solve randomly generated instances with up to 15 nodes with gaps about 10%. Considering a relaxation (based on cut inequalities) of a network design problem for one basic technology Barahona [3] solves instances with up to 47 nodes with an accuracy of 5-10%.

Discrete capacities and diversification/reservation

Dahl and Stoer [6, 7] were the first to consider a discrete capacity structure and survivability issues in the same model (for *Norwegian Telecom Research*). They solve a large number of the instances (from 37 to 118 nodes with a very sparse supply graph) to optimality. But they also report considerable difficulties with some of the instances. Their work is the basis of our investigations. We are, in particular, grateful for the contributions of Mechthild Stoer who was a member of the project team at ZIB (Konrad-Zuse-Zentrum für Informationstechnik Berlin) in the beginning of our study. Models that incorporate discrete capacities and survivability are the main topic of the following sections.

2 DISCNET – Models

After having reviewed various models considered by other authors, we now present the models that we developed for e-plus Mobilfunk GmbH, one of the mobile-communication service providers in Germany. We distinguish between different capacity and survivability models. In particular, we consider two ways to model the **discrete capacity structure** and three ways to achieve **survivability** in the network. Any combination of the two leads to a different mixed-integer programming formulation. All of these are integrated in our network dimensioning tool DISCNET (DImensioning Survivable Cellular-phone NETworks). Our solution approach is sketched in Section 3 and computational results are described in Section 4.

2.1 Capacity Models

We distinguish between two different capacity models. First, we consider the case of an arbitrary, but finite, set of possible capacities for each edge of the supply graph. Then we discuss the case where a small set of basic capacities is given such that each capacity is an integral multiple of each smaller capacity, a reasonable assumption in telecommunications. In Germany, for instance, *Deutsche Telekom* offers the basic capacities 30, 480 and 1920 channels and in the United States the technologies DS0, DS1 and DS3 (DS = Digital Signal Level) come along with capacities 1, 24 and 672.

In our practical application, it was natural to assume that every edge $e \in E$ of the supply graph is already equipped with an initial capacity $C_e^0 \in \mathbb{Z}_+$ (possibly $C_e^0 = 0$) of cost $K_e^0 = 0$, the so-called **free capacity**. This applies to both capacity models.

DISCRETE CAPACITIES

For each $e \in E$, there is a finite set of capacities specified by the following data:

- $T_e \in \mathbb{Z}_+$, where T_e is the **number of possible additional capacities** that can be installed,
- $C_e^t \in \mathbb{Z}_+$, $1 \leq t \leq T_e$, the **potential capacities** (we assume that $C_e^0 < C_e^1 < \cdots < C_e^{T_e}$),
- $K_e^t \in \mathbb{Q}_+$, $1 \leq t \leq T_e$, the **cost** of installing capacity C_e^t .

It has turned out to be useful to call the capacities $C_e^1, \ldots, C_e^{T_e}$ **breakpoint capacities**, and hence T_e the number of **breakpoints**, and to consider the **incremental capacities and costs**

- $c_e^t := C_e^t - C_e^{t-1}$, $1 \leq t \leq T_e$,
- $k_e^t := K_e^t - K_e^{t-1}$, $1 \leq t \leq T_e$,

instead of the original values. For notational reasons, we set $c_e^0 := C_e^0$ and $k_e^0 := K_e^0$.
For each edge $e \in E$ we introduce an ordered set of **0/1 variables** $x_e^0 \geq x_e^1 \geq \cdots \geq x_e^{T_e}$. Since we assume that a free capacity C_e^0 is always installed, we set $x_e^0 := 1$. Choosing capacity C_e^τ , $0 \leq \tau \leq T_e$, is equivalent to setting $x_e^0 = x_e^1 = \cdots = x_e^\tau = 1$ and $x_e^{\tau+1} = \cdots = x_e^{T_e} = 0$.
The objective is to minimize the total cost of installing the necessary capacities on the edges of the supply graph. This is formulated as

$$\min \sum_{e \in E} \sum_{t=1}^{T_e} k_e^t x_e^t \,. \tag{1}$$

The 0/1-variables associated with a supply edge must satisfy the **ordering constraints**

$$1 = x_e^0 \geq x_e^1 \geq \cdots \geq x_e^{T_e} \geq 0 \qquad \text{for all } e \in E \,, \tag{2}$$

and the **integrality constraints**

$$x_e^t \in \{0,1\} \qquad \text{for all } e \in E \text{ and } t = 1, \ldots, T_e \,, \tag{3}$$

by definition. For notational convenience, we introduce **auxiliary variables**

$$y_e := \sum_{t=0}^{T_e} c_e^t x_e^t \qquad \text{for all } e \in E \,, \tag{4}$$

representing the capacities installed on the supply edges.

Divisible Basic Capacities

In the second capacity model we can install on all supply edges combinations of a common set of basic capacities. These basic capacities satisfy the property that each capacity is an integral multiple of each smaller capacity.

We are given a set $T = \{\tau_1, \dots, \tau_n\}$ of **technologies**, one for each different type of line that can be installed on a supply edge. With each technology $\tau \in T$, we associate a **basic capacity** C^τ (assuming w.l.o.g. $C^{\tau_1} \leq \dots \leq C^{\tau_n}$) and the edge dependent installation **costs** (which include a fixed cost and a length-dependent cost that varies with the total length of a link). We assume that the basic capacities satisfy the divisibility property, i.e.,

$$\frac{C^{\tau_{i+1}}}{C^{\tau_i}} \in \mathbb{Z}_+$$

for all $i = 1, \dots, n-1$. We refer to the smallest basic capacity as the **unit capacity**.

For each supply edge $e \in E$ we are given a set $t(e) \subseteq T$ of available technologies. The capacities that can be installed on edge e are integer combinations of the basic capacities of the available technologies. For this purpose we introduce, for every supply edge $e \in E$ and every technology $\tau \in t(e)$, a nonnegative integer variable x_e^τ to denote the integral multiple of C^τ. The variables x_e^τ may be restricted by an upper bound u_e^τ. For each $\tau \in t(e)$ we denote by K_e^τ the cost of installing the basic capacity C^τ on supply edge $e \in E$.

Again, the objective is to minimize the total cost of installing the necessary capacities on the edges of the supply graph. This is formulated as

$$\min \sum_{e \in E} \sum_{\tau \in t(e)} K_e^\tau x_e^\tau \,. \tag{5}$$

The constraints that must be satisfied are the nonnegativity, the integrality, and possibly the upper bound constraints

$$0 \leq x_e^\tau \leq u_e^\tau \qquad \text{and} \qquad x_e^\tau \in \mathbb{Z}_+ \qquad \text{for all } e \in E \text{ and all } \tau \in t(e) \,, \tag{6}$$

where the capacity y_e of a supply edge $e \in E$ is

$$y_e = C_e^0 + \sum_{\tau \in t(e)} C^\tau x_e^\tau. \tag{7}$$

2.2 Combining Capacities, Demands and Routings

Combining any of the two capacity models with the multicommodity flow conditions for the non-failure situations, to be described below, we obtain the basic mixed-integer programming formulations mentioned in the beginning.

For the network we want to design, we also wish to determine the routings of the demands $uv \in D$ for each **operating state** s of the network. The operating states are

- the **normal** state ($s = 0$), which is the state with all nodes and edges operational, and, in case we wish to consider failure situations,

- the **failure** states, which (in our case) are the states with a single node u ($s = u$) or a single edge e ($s = e$) nonoperational.

We denote by $G_s = (V_s, E_s)$ the supply graph for the operating state s, where V_s is the set of nodes that are still operational in operating state s, and, likewise, E_s is the set of the operational edges in operating state s. Similar notational conventions apply to the demand graph.

We impose another condition to the paths used to route the demands in the normal operating state. For each demand $uv \in D$ we introduce

- $\ell_{uv} \in \mathbb{Z}_+$, the **path length restriction**; ℓ_{uv} is the maximum number of supply edges allowed in any path on which demand between u and v is routed.

For each operating state s and each demand edge $uv \in D_s$, let $\mathcal{P}(s, uv)$ denote the set of **valid** $[u, v]$**-paths** in G_s. If s is the normal operating state, a $[u, v]$-path in $G = G_0$ is valid if its length (number of edges) is at most ℓ_{uv}. We call such a path **short**. If s is a failure state then any $[u.v]$-path in G_s is valid. For each operating state s, each edge $uv \in D_s$, and each path $P \in \mathcal{P}(s, uv)$, we define a variable $f(s, uv, P)$, called **flow** or **routing variable**, that represents the communication traffic between the nodes u and v routed on path P in operating state s.

The constraints for the routings in the normal operating state are the capacity, demand and nonnegativity constraints. The capacity constraints imply that for each supply edge $e \in E$ the flow through e may not exceed its capacity; the routing variables must be chosen in such a way that all the demands d_{uv} in the normal operating state are satisfied. This yields the demand constraints. Putting this together with the nonnegativity of the routing variables we get

$$\sum_{uv \in D} \sum_{P \in \mathcal{P}(0,uv): e \in P} f(0, uv, P) \leq y_e \quad \text{for all } e \in E\,, \tag{8}$$

$$\sum_{P \in \mathcal{P}(0,uv)} f(0, uv, P) = d_{uv} \quad \text{for all } uv \in D\,, \tag{9}$$

$$f(0, uv, P) \geq 0 \quad \text{for all } uv \in D \text{ and } P \in \mathcal{P}(0, uv)\,. \tag{10}$$

Combining the capacity constraints (1) - (4) or (5) - (7) with the routing constraints (8), (9), (10) for the normal operating state we obtain two network design models that do not incorporate survivability at all:

(ND1) defined by the DISCRETE CAPACITIES constraints (1) - (4) and (8), (9), (10), and
(ND2) defined by the DIVISIBLE BASIC CAPACITIES constraints (5) - (7) and (8), (9), (10)

If we only use one or two different technologies and if we do not impose any length restriction on the paths our basic model (ND2) is exactly the model studied in [12, 5, 3] (see the paragraph "Discrete capacities and no survivability" in Section 1).

2.3 Survivability Models

We now consider three different concepts to model survivability. Using diversified paths to satisfy the demands, i.e., node-disjoint paths that do not carry more than a certain percentage of a given demand, it is possible to achieve survivability with the advantage of no maintenance in case of a component failure. However, as an obvious disadvantage it is not possible to guarantee saving 100% of the demands in case of a failure; and, e.g., for more than 50% of each demand to survive, at least three paths are required for each demand. An additional empirical drawback is the high cost of implementing diversification. In order to satisfy a specified percentage of each demand, in case a network component fails, the two other concepts make use of rerouting. If the network designer opts for the reservation concept he also decides that, in case of a component failure, the routings for all demands can (and in general have to be) changed. An alternative is to consider the rerouting of those demands only that are affected by a failure. Both concepts require higher maintenance efforts, but yield smaller network dimensioning costs.

Diversification

To implement this concept, the network designer has to specify

- the **diversification** parameter δ_{uv}, $0 < \delta_{uv} \leq 1$, for all $uv \in D$; δ_{uv} is the maximum fraction of the demand d_{uv} allowed to flow through any supply edge or node (other than nodes u and v).

The **node-flow constraints**

$$\sum_{P \in \mathcal{P}(0,uv): w \in P} f(0,uv,P) \leq \delta_{uv} d_{uv} \qquad \text{for all } uv \in D \text{ and } w \in V \backslash \{u,v\}\ , \tag{11}$$

and the **edge-flow constraints**

$$f(0,uv,P) \leq \delta_{uv} d_{uv} \qquad \text{for all } uv \in D \text{ and } P = \{uv\}\ , \tag{12}$$

are the **diversification constraints.** The summation in the node-flow constraints is over all short paths between nodes u and v that contain node w. These constraints restrict the amount of flow dedicated to a particular demand that goes through a particular node, i.e., they ensure that in the normal operating state, no more than a fraction δ_{uv} of the total demand d_{uv} between nodes u and v flows through a single node w. The node-flow constraints imply that every edge $e \in E$ carries no more than $\delta_{uv} d_{uv}$ of the traffic between u and v, unless $e = uv$. To cover the latter case, the edge-flow constraints are used. These are employed only, of course, if E contains edges between u and v (which are considered as paths $P = \{uv\}$). The constraints (11) and (12) yield that the flow between u and v is **diversified**, i.e., is routed on at least $\lceil \frac{1}{\delta_{uv}} \rceil$ node-disjoint paths.

Reservation

Here, the network designer has to specify

- the **reservation** parameter ρ_{uv}, $0 \leq \rho_{uv} \leq 1$, for all $uv \in D$; ρ_{uv} is the fraction of the demand d_{uv} that must be satisfied in a single node failure or a single supply edge failure.

The **capacity** and **demand constraints** for failure situations are

$$\sum_{uv \in D_s} \sum_{P \in \mathcal{P}(s,uv): e \in P} f(s, uv, P) \leq y_e \quad \text{for all } s \neq 0 \text{ and } e \in E_s, \tag{13}$$

$$\sum_{P \in \mathcal{P}(s,uv)} f(s, uv, P) = \rho_{uv} d_{uv} \quad \text{for all } s \neq 0 \text{ and } uv \in D_s\,. \tag{14}$$

With inequality (13), the flow through a supply edge does not exceed its capacity in failure situations, and with inequality (14), $100\rho_{uv}$ percent of the demand d_{uv} survive the failure of an edge or a node, in case the network routing is switched to the new routings.

Rerouting of Affected Demands

For this concept the network designer has to specify

- the **rerouting** parameter σ_{uv}, $0 \leq \sigma_{uv} \leq 1$ for all $uv \in D$; σ_{uv} is the fraction of the demand d_{uv} that must be satisfied in a single node failure or a single supply edge failure, **without** rerouting the paths that are not affected by the particular failure situation.

Here, the **capacity** and **demand constraints** for failure situations have to take care of the unaffected normal operating state routings. Since we possibly impose a length-restriction on the paths used to route a demand in the normal operating state, we introduce for each demand $uv \in D_s$ and each failure situation s ($\neq 0$) the symbol $\mathcal{P}_s(0, uv)$ to denote the set of short paths in $\mathcal{P}(s, uv)$, i.e., $P \in \mathcal{P}_s(0, uv)$ iff $P \in \mathcal{P}(s, uv)$ and P has at most l_{uv} supply edges. We get the additional constraints

$$\sum_{uv \in D_s} \Big(\sum_{P \in \mathcal{P}_s(0,uv): e \in P} f(0, uv, P) + \sum_{P \in \mathcal{P}(s,uv): e \in P} f(s, uv, P) \Big) \leq y_e \quad s \neq 0,\ e \in E_s, \tag{15}$$

$$\sum_{P \in \mathcal{P}_s(0,uv)} f(0, uv, P) + \sum_{P \in \mathcal{P}(s,uv)} f(s, uv, P) \geq \sigma_{uv} d_{uv} \quad s \neq 0,\ uv \in D_s. \tag{16}$$

Note, in case the normal operating state routings suffice to satisfy a demand $uv \in D$ in a particular failure situation s, i.e., if $\sum_{P \in \mathcal{P}_s(0,uv)} f(0, uv, P) \geq \sigma_{uv} d_{uv}$, there is no rerouting at all necessary. Thus the maintenance effort is reduced.

2.4 Mixed-Integer Linear Programming Formulations

Combining the two capacity and the three survivability models we get six mixed-integer linear programming models for capacitated network dimensioning problems taking survivability issues into account.

1. Divisible basic capacities + Diversification
2. Divisible basic capacities + Reservation
3. Divisible basic capacities + Rerouting of affected demands
4. Discrete (arbitrary) capacities + Diversification
5. Discrete (arbitrary) capacities + Reservation
6. Discrete (arbitrary) capacities + Rerouting of affected demands

Dahl and Stoer investigate the models 4. and 5. in [6, 7]. If we consider continuous instead of discrete capacities then 2. (or 5.) is the model Minoux investigates in [13], and 3. (or 6.) is the model of Lisser, Sarkissian and Vial (see [11]), if the affected demands are rerouted in a separate network.

3 Algorithmic Approach

In this section we provide a high-level description of the cutting-plane algorithm, based on linear programming and polyhedral combinatorics, we developed to solve the problems described in the previous section. Figure 1 shows the flow chart of the algorithm.

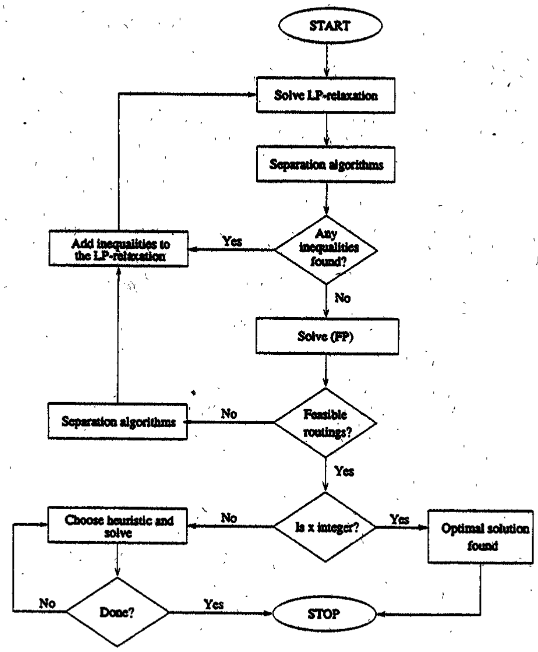

Figure 1: Flow chart of the algorithm.

The algorithm consists of three main parts:

1. the multicommodity flow problems (FP) to test the feasibility of given capacities,
2. the cutting plane part to obtain a lower bound for the optimal solution value, and,
3. the heuristic algorithms to obtain feasible solutions.

The multicommodity flow problems are given by the constraints (8), (9), (10) plus the appropriate survivability constraints. These are, for *diversification*, the constraints (11) and (12); for *reservation* the constraints (13) and (14), and for the survivability model *rerouting of affected demands* the constraints (15) and (16). We apply a solution approach suggested by Minoux and already used by Dahl and Stoer. See [13], [6] and [1] for details. If we use *rerouting of affected demands* as the survivability model then the multicommodity flow problems are much more complicated than in the other two cases, because the different operating states are linked through the normal operating state routing.

In the cutting plane part we solve at each iteration a linear programming relaxation of the integer program defined by the decision variables x of the chosen capacity model. In particular, we relax the integrality conditions and we add iteratively inequalities valid for the polyhedron defined by the integral feasible x-vectors.

For the capacity model DISCRETE CAPACITIES, the valid inequalities we use are

1. strengthened band inequalities, introduced by Dahl and Stoer [6],
2. strengthened metric inequalities, see [1], and
3. diversification bands, see [1], if *diversification* is the chosen survivability model,

and, for the capacity model DIVISIBLE BASIC CAPACITIES, the valid inequalities we use are

1. partition inequalities introduced by Pochet and Wolsey [16],
2. strengthened partition inequalities, see [2],
3. strengthened metric inequalities, see [2], and
4. diversification-cut inequalities, see [2], if *diversification* is the chosen survivability model.

Various exact and heuristic separation algorithms (see [6, 16, 1, 2]) have been developed for these classes of inequalities. We iterate the separation algorithms and the multicommodity-flow algorithms, for the solution of (FP), as long as we can find violated inequalities. When this process stops, the fractional solution in x-variables permits feasible solutions to all multicommodity-flow problems, i.e., given these fractional capacities we can determine routings for all demands in all operating states. If the x-variables are integral we have found an optimal solution. Otherwise, we resort to heuristic algorithms (see [1, 2]) to obtain "good" integer solutions. In our present implementation we run 12 heuristics which are parameter controlled variants of two different design principles.

The cutting plane phase provides a lower bound and the best heuristic solution provides an upper bound for the optimal solution value (the minimum cost for dimensioning the network). Thus, our solutions have a guaranteed quality.

4 Computational Results

In this paper we have indicated various ways to define and mathematically model "survivability" of telecommunication networks. We have described a general model that integrates several concepts and sketched a cutting plane algorithm to solve it. This algorithm is the core of the network design tool DISCNET that we developed for e-plus Mobilfunk GmbH. The tool has been implemented in C++ and is in use at e-plus.

Our algorithm is able to produce reasonable solutions for practical instances in acceptable running times. The e-plus instance sizes range from 11 nodes, 34 supply edges, 24 demand edges to 17 nodes, 64 supply edges, 106 demand edges resulting in mixed-integer programs with up to 1,000 integral and 100,000 continuous variables. The time to compute a lower and an upper bound for the optimum value with our cutting plane algorithm and our set of 12 heuristics, respectively, ranges from a few seconds up to several hours, depending on the used survivability model. The solutions calculated with *diversification* as survivability model turned out to be the most expensive. We expected this result since, in this model, the edge capacities for the normal operating state must be chosen in such a way that all failure situations can be handled without rerouting of traffic. The cheapest solutions are produced with the *reservation* concept, which, however, requires additional hard- and software for the rerouting effort in failure cases. The best compromise between cost and maintenance effort seems to be obtainable with the survivability model *rerouting of affected demands.* The solutions are not much more expensive than the comparable *reservation* solutions, and the rerouting effort is relatively small.

The quality of the solutions our methodology produces is, from a theoretical point of view, not satisfactory yet. The "integrality gap" may be small for sparse graphs with few possible capacities. But upper and lower bound may differ by 50% – 60% for particularly difficult cases. We believe that this is due to a poor lower bound since, in this complex mix of models, our cutting planes tend to be rather weak. Further investigations of the polyhedral structure of the convex hull of the feasible solutions are necessary.

However, the solutions we determine are about 15% – 20% better than the ones produced by the network designers (using more traditional techniques) and result in considerable savings. Moreover, with our approach we do guarantee that all side constraints are satisfied. Our program DISCNET has become an intensively used planning tool. The designers compute the network topology and edge capacities under several different assumptions and parameter settings and, after analyzing all results, choose a network dimensioning suitable for the company needs that provides a good compromise between installation and maintenance costs.

References

[1] D. Alevras, M. Grötschel, and R. Wessäly. A network dimensioning tool. Preprint SC 96-49, Konrad-Zuse-Zentrum für Informationstechnik, Berlin, 1996.

[2] D. Alevras, M. Grötschel, and R. Wessäly. Cost–efficient network synthesis from leased lines. Preprint SC-97-22, Konrad-Zuse-Zentrum für Informationstechnik, Berlin, 1997.

[3] F. Barahona. Network design using cut inequalities. *SIAM Journal on Optimization*, 6(3):823–837, 1996.

[4] D. Bienstock, S. Chopra, O. Günlück, and C.-Y. Tsai. Minimum cost capacity installation for multicommodity network flows. *Mathematical Programming (Special Issue on Computational Integer Programming)*, (to appear), 1995.

[5] D. Bienstock and O. Günlük. Computational experience with a difficult mixed-integer multicommodity flow problem. *Mathematical Programming*, 68:213–237, 1995.

[6] G. Dahl and M. Stoer. MULTISUN – Mathematical model and algorithms. Technical Report TF R 46/92, Televerkets Forskningsinstitutt, 1992.

[7] G. Dahl and M. Stoer. A polyhedral approach to multicommodity survivable network design. *Numerische Mathematik*, 68:149–167, 1994.

[8] M. Grötschel, C.L. Monma, and M. Stoer. Computational results with a cutting plane algorithm for designing communication networks with low-connectivity constraints. *Operations Research*, 40(2):309–330, 1992.

[9] M. Grötschel, C.L. Monma, and M. Stoer. Facets for polyhedra arising in the design of communication networks with low-connectivity constraints. *SIAM Journal on Optimization*, 2(3):474–504, 1992.

[10] M. Grötschel, C.L. Monma, and M. Stoer. *Design of survivable networks*, volume Network Models of *Handbooks in Operations Research and Management Science*, chapter 10, pages 617–672. North-Holland, 1995.

[11] A. Lisser, R. Sarkissian, and J.P. Vial. Optimal joint syntheses of base and spare telecommunication networks. Technical report, University of Genève, November 1995.

[12] T.L. Magnanti, P. Mirchandani, and R. Vachani. Modeling and solving the two–facility capacitated network loading problem. *Operations Research*, 43(1):142–156, 1995.

[13] M. Minoux. Optimum synthesis of a network with non-simultaneous multicommodity flow requirements. In P. Hansen, editor, *Studies on Graphs and Discrete Programming*, pages 269–277. North–Holland Publishing Company, 1981.

[14] C.L. Monma, B.S. Munson, and W.R. Pulleyblank. Minimum-weight two-connected spanning networks. *Mathematical Programming*, 46:153–171, 1990.

[15] C.L. Monma and D.F. Shallcross. Methods for designing communications networks with certain two-connected survivability constraints. *Operations Research*, 37(4):531–541, 1989.

[16] Y. Pochet and L.A. Wolsey. Network design with divisible capacities: Aggregated flow and knapsack formulation. In *Proceedings IPCO2*, pages 150–164, Pittsburgh, 1992. Carnegie Mellon University.

[17] M. Stoer. *Design of Survivable Networks*. Lecture Notes in Mathematics. Springer, 1992.