

# Survivable Mobile Phone Network Architectures: Models and Solution Methods

Dimitris Alevras, IBM

Martin Grötschel, Konrad-Zuse-Zentrum für Informationstechnik

Peter Jonas, E-Plus Mobilfunk

Ulrich Paul, o.tel.o

Roland Wessäly, Konrad-Zuse-Zentrum für Informationstechnik

**ABSTRACT** In this article we present a mixed-integer programming model for the problem of designing a survivable capacitated network; and describe a cutting plane algorithm for its solution. The model and the solution methods are integrated in our network dimensioning tool, DISCNET. Given a communication demand between each pair of switching nodes in a region, the task is to determine the topology of a telecommunication network connecting the given nodes and to select, from a given set of valid values, a capacity for each potential physical link such that the communication demands are satisfied, even if a network component fails. A solution consists of the chosen links and their capacity, as well as the routings for each demand, in the case of failure-free operation and the case of single component (node or link) failure. We suggest two alternative models to deal with failures of single network components. The first employs diversified paths to guarantee the routing of a specified fraction of each demand without rerouting effort; the second allows rerouting in failure situations. At the end we discuss alternative ways to implement survivability using these two models.

ality which describes the behavior of the network. The switching level takes control of the grade of service (GoS), which normally is defined as the number of blocked (lost) calls in the network. The planning result of this level, which takes into account the number of switching nodes and the selected routing schemes, will be used as input for the transport level. Using multiplexing systems is key to optimizing the allocation of physical resources when mapping the demand from the switching level to physical resources. This is because economies of scale are realized by multiplexing 64 kb/s or even less (16 or 8 kb/s) channels

The design, dimensioning, and administration of survivable telecommunications networks (i.e., networks that survive the failure of certain components) are becoming more and more important. This is because overall service quality has become a major competition criterion for telecommunications services. End-to-end survivability is not only a subject of broadband networks, as in the EC-sponsored IMMUNE project, which is part of the RACE program, but also for smaller telecommunications networks such as mobile networks, where overall service quality is eminent and vital in a highly competitive environment. However, the right balance between costs and quality must be determined by the design engineers.

Certain protection mechanisms have been developed and applied in SDH technology and in digital cross-connect (DCC) systems, that is, synchronous digital hierarchy (SDH) or plesiochronous digital hierarchy (PDH) technology. Diverse routed protection, 1+1 automatic protection rings (APR), path protected switched routing (PPS), multiplexer section protected rings (MSP rings), and restoration at the DCC level are a few among others. Together with component redundancy and dynamic restoration methods, these are applied to help specific parts of specific networks to survive failures of one or more of their components.

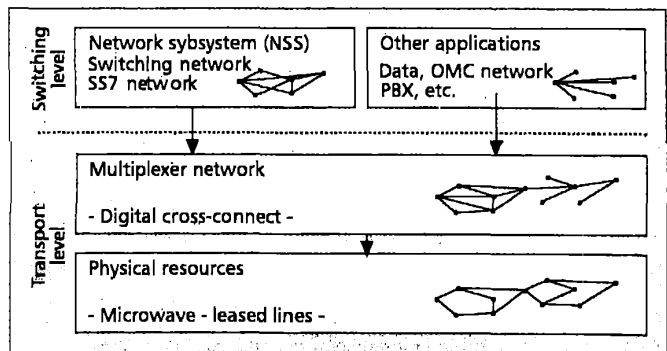
What is still missing is an integrated approach to the network design problem where, at the same time, cost effectiveness, survivability, and network management aspects are taken into account to achieve a solution that appears economically efficient from various points of view.

Figure 1 shows the architecture of a typical mobile telecommunications network, including a multiplexer network level. Each level is equipped with a certain function-

for different applications and services (mobile communication, office data communication, corporate networks, etc.). Using SDH technology, the lowest multiplexing level is 2.048 Mb/s, which will be mapped into virtual containers. Quality of service (QoS), in terms of availability of used physical transmission systems (e.g., leased lines or microwave), is considered on this level.

Survivability is considered during the design process by using the mentioned protection mechanisms at the transport level. Survivability in this context is the fraction of the demand that is satisfiable in a failure case (e.g., if a physical link or node fails).

In this article, we consider the problem of designing a survivable telecommunications network, that is, the problem of selecting from a discrete set of capacities which one to install on each link of the physical network and deciding how to route each



■ Figure 1. Typical mobile telecommunications network architecture.

demand (even in the case of a single node or single edge failure) at minimum installation cost. Additional restrictions to the percentage of a demand routed through a particular node or edge of the network, and the length of the paths between two demand nodes are considered. We model the problem as a mixed-integer programming problem and present a cutting plane algorithm to solve it. Due to the complexity of the problem, we do not expect to get optimal solutions. Instead, we get low-cost solutions with a quality guarantee, which is an upper bound to the gap between the solution value and the (unknown) optimum.

Variations of our problem can be found in the literature. Most of these models consider either nondiscrete capacities and survivability (e.g., [1, 2]), or they consider discrete capacities without survivability (e.g., [3-5]). These later studies, however, restrict the possible capacities to multiples of two base-capacities. Dahl and Stoer [6, 7] study a problem similar to ours but without imposing length restrictions on the paths between demand nodes.

Here we present a network dimensioning tool, which we call DISCNET (Dimensioning of Survivable Cellular phone NETWORKS), which was developed for E-Plus Mobilfunk GmbH, Germany. The rest of the article is organized as follows. In the next section we formally define the problem and present the model. A high-level description of the algorithm is given in the third section, while in the fourth section we describe the input data and present typical output of the tool. In the fifth section we discuss three methods of setting parameters of the model to realize survivability in the network.

## FORMULATION OF THE PROBLEM

Formally, the problem is defined as follows. The input consists of two graphs, the *supply graph*  $G = (V, E)$  and the *demand graph*  $H = (V, D)$ . Both graphs share the same node-set  $V$ . The edge set  $E$  of the supply graph consists of all the possible links to be considered. Different types of links (e.g., microwave, leased lines) are modeled using *parallel edges*. An edge between two nodes in the demand graph denotes an existing demand between the two nodes.

With each edge  $e \in E$  of the supply graph we associate a *cost type* that provides information about the possible capacities available for the particular link, as well as the related costs. We refer to each possible capacity for a particular link as a *breakpoint capacity*. The number  $T_e$  of possible capacities is also referred to as the *number of breakpoints*. The cost type also includes a "free" capacity which is the capacity already installed on a link (and comes at no cost), and the incremental capacity  $m_e^t$  and incremental cost  $k_e^t$  for each breakpoint  $t$  ( $0 \leq t \leq T_e$ ), which correspond to the additional capacity and additional cost, respectively, resulting from going from breakpoint  $t-1$  to breakpoint  $t$ .

With each edge  $uv \in D$  of the demand graph, we associate four parameters:

- $d_{uv}$  which is the demand between nodes  $u$  and  $v$
- $\delta_{uv}$  which is the *diversification* parameter, that is, the maximum fraction of the demand  $d_{uv}$  allowed to flow through any edge or node (other than nodes  $u$  and  $v$ ) (Fig. 2)

- $r_{uv}$  which is the *reservation* parameter, that is, the minimum fraction of the demand  $d_{uv}$  which must be satisfied in a single node or a single edge failure (Fig. 2)
- $l_{uv}$  which is the *path length restriction*, that is, the maximum number of edges allowed in any path on which the demand  $d_{uv}$  is routed

All these parameters correspond to restrictions that a *feasible network* must satisfy.

Eventually, we wish to determine the capacity to install on each link (edge of the supply graph) to have a low-cost or minimal-cost feasible network. In addition, we wish to have the routings of the demands for each *operating state* of the network. We use the index  $s$  to denote an operating state. The operating states of the network are:

- The *normal* state ( $s = 0$ ), which is the state with all nodes and all edges operational
- The *failure* states, which 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  $E_s$  is the set of the operational edges in operating state  $s$ . Similar notational conventions apply to the demand graph.

We model the problem as a *mixed-integer linear programming problem*. There are two types of variables in our model, the *decision* and the *routing* variables.

The decision variables are zero-one variables indicating which capacity is chosen for the edges of the supply graph. For each  $e \in E$  the decision variable  $x_e^t$  equals 1 for all breakpoints  $t \leq \tau$ , where  $\tau \leq T_e$  is the chosen breakpoint, and equals 0 otherwise. The routing variables are needed in order to obtain the routings in each operating state. The routing variable  $f(uv, s, P)$ , for an operating state  $s$  and a demand edge  $uv \in D_s$  denotes the part of the demand  $d_{uv}(s=0)$  or  $r_{uv}d_{uv}$  (else) that is routed on the *short path*  $P \in P(uv, s)$ , where  $P(uv, s)$  is the set of all short paths in  $G_s$ . By short path, here we mean any path that satisfies the path length restriction  $l_{uv}$ .

The objective is to minimize the total capacity installation cost. With this information we can write the mixed-integer linear programming model as follows:

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

subject to

$$0 \leq x_e^{T_e} \leq \dots \leq x_e^1 \leq x_e^0 = 1 \quad \forall e \in E \quad (1)$$

$$x_e^t \in \{0, 1\} \quad \forall e \in E, \forall t = 1, \dots, T_e \quad (2)$$

$$y_e = \sum_{t=0}^{T_e} m_e^t x_e^t \quad \forall e \in E \quad (3)$$

$$\sum_{uv \in D_s} \sum_{P \in P(uv, s); e \in P} f(uv, s, P) \leq y_e \quad \forall s \quad \forall e \in E_s \quad (4)$$

$$\sum_{P \in P(uv, 0)} f(uv, 0, P) = d_{uv} \quad \forall uv \in D \quad (5)$$

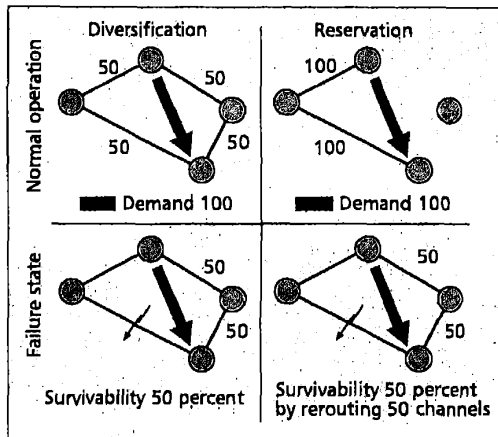
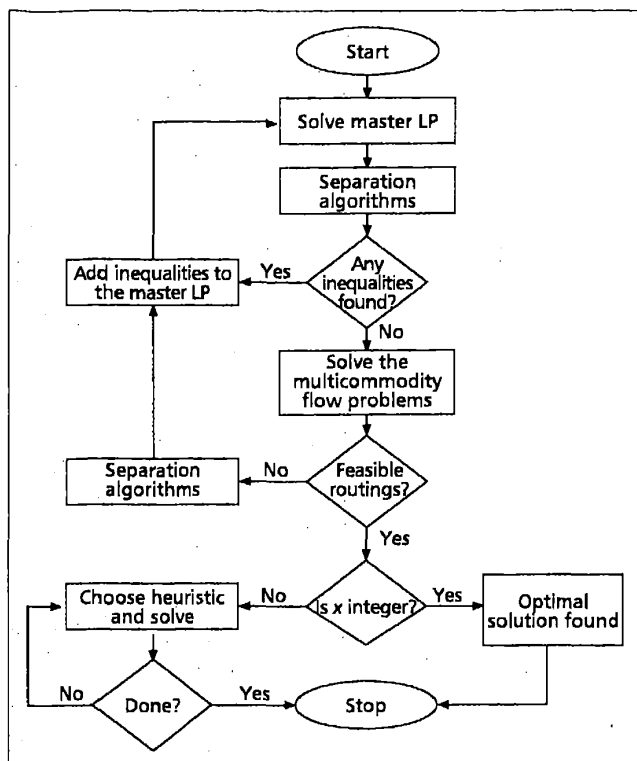


Figure 2. Diversification and reservation.



■ Figure 3. Flow chart of the algorithm.

$$\sum_{P \in P(uv,s)} f(uv,s,P) = r_{uv} d_{uv} \quad \forall s \neq 0 \quad \forall uv \in D_s \quad (6)$$

$$\sum_{P \in P(uv,0); w \in P} f(uv,0,P) \leq \delta_{uv} d_{uv} \quad \forall w \in V \quad w \neq u, v \quad (7)$$

$$f(uv,0,P) \leq \delta_{uv} d_{uv} \quad \forall uv \in D, P = \{uv\} \quad (8)$$

$$f(uv,s) \geq 0 \quad \forall s, \forall uv \in D_s, \quad \forall P \in P(uv,s) \quad (9)$$

A feasible solution to our problem is an  $x$  vector that has integer (zero-one) components and that corresponds to a feasible capacity vector  $y$ . The capacity vector  $y$ , which is calculated from the  $x$  vector from Eq. 3, is feasible if it permits feasible routings for all operating states, that is, if the system of linear equalities and inequalities, Eqs. 4-9, has a feasible solution for the given  $y$ .

### THE ALGORITHMIC APPROACH

In this section we give a high-level description of a cutting-plane algorithm that we developed to solve the problem described in the previous section; see [8, 9] for a general description of a cutting plane algorithm. This type of algorithm has been applied successfully to other problems in telecommunications [10-12].

In Fig. 3 we show the flow chart of the algorithm. The algorithm consists of three main parts:

- The cutting plane part, which deter-

- mines a lower bound for the objective function value and a starting point for the LP-based heuristics
- Multicommodity flow problems, which are solved to determine whether a given set of capacities permits feasible routings or not
- The heuristic algorithms used to obtain feasible solutions


In the cutting plane part we solve the *master LP*, which contains constraints in  $x$  variables only. In particular, these are the ordering constraints of Eq. 1 and a subset of the valid inequalities for the 0-1 polytope of feasible  $x$  vectors. The valid inequalities (cutting planes) we use are:

- Strengthened band inequalities, introduced by Dahl and Stoer [6]
- Strengthened metric inequalities [13]
- Diversification cuts [13]

These inequalities are identified automatically during the execution of the algorithm via so-called separation routines, which are algorithms that, given an  $x$  vector, identify a violated inequality if one exists. There are two reasons a separation algorithm might fail to identify a violated inequality: either such an inequality does not exist, or the separation algorithm is heuristic, that is, it does not guarantee that a violated inequality will be found even if it exists. If we fail to separate a violated inequality we switch to the multicommodity flow problems to determine whether the (possibly fractional) solution obtained for the master LP corresponds to a feasible capacity vector  $y$ . If so, and if the  $x$  solution from the master LP is integer, we have found an optimal solution and we are done. If the feasibility test fails for at least one operating state, then a violated metric inequality [14, 15] has been identified. From the derived metric inequalities we try to separate violated strengthened band inequalities. The underlying problem for the separation is the multiple-choice knapsack problem [16]. We apply a dynamic-programming-based exact procedure and a heuristic procedure suggested by Stoer and Dahl [6] to solve the separation problem for strengthened band inequalities. If we fail to generate violated strengthened-band inequalities, we derive strengthened metric inequalities, applying a divide-and-round procedure to the metric inequalities. The identified violated inequalities are then added to the master LP and the whole procedure is repeated. If feasible routings are found for all operating states and we cannot find violated inequalities, but the  $x$  variables are not integer, we resort to various LP-based heuristics to obtain "good" integer solutions. The heuristics are of two types. In the first type one branch of the branch-and-cut tree (see, e.g., [9]) is examined. Different selection criteria for the branching variable  $x_e^i$  give rise to different heuristics of this type. In the other type, we first derive a feasible capacity vector from the fractional master LP solution; then we try to improve the solution using various criteria to reduce the capacities of the supply edges.

The cutting plane phase provides a lower bound  $z_{LP}$ , and the best heuristic solution provides an upper bound  $z_{IP}$ . Thus, we get a quality guarantee for the best solution found by the algorithm, that is, an upper bound to the gap between the best solution found and an optimal solution, given by the quantity

$$\frac{z_{IP} - z_{LP}}{z_{LP}} 100.$$



■ Figure 4. Supply graph showing all possible transmission systems.

## PLANNING PROCESS AND RESULTS


In this section we describe how the supply and demand graphs are obtained, and present typical output of the algorithm.

In principle, the supply graph can be a complete graph with parallel edges. However, it is evident from the model of the second section that careful selection of the possible links and breakpoint capacities is desired, since this will reduce the number of integer variables in the model. Moreover, we have observed in practice that working with complete graphs significantly increases the running time of the algorithm and, in most cases, gives solutions the quality of which is not better than that of solutions derived with supply graphs with fewer edges. We remark that the supply graph must satisfy certain connectivity requirements that are imposed by the diversification and reservation parameters. In Fig. 4 we show an example of a supply graph.

The demands between nodes (i.e., the edges of the demand graph) are obtained in the switching planning process by applying a *call*-based routing algorithm to map the forecast traffic between nodes (in Erlangs) to logical demands (in channels). The other input parameters of the demand graph are provided as follows. The length restriction can be any integer number bigger than 1. The reservation parameter  $r_{uv}$  is set to a value between 0 and 1 (i.e.,  $0 \leq r_{uv} \leq 1$ ), while the diversification parameter  $\delta_{uv}$  is set to a positive value between 0 and 1 (i.e.,  $0 < \delta_{uv} \leq 1$ ). Figure 5 shows an example of a demand graph.

The primary result is the dimensioning of the physical network, that is, the assignment of capacities to the links which permit the satisfaction of the demands in the normal state and of a given percentage of the demands in each single failure state (Figs. 6 and 7 for two examples of the output for the graphs of Figs. 4 and 5). Depending on the applied heuristics we get different network topologies with different total costs. Thus, the planner can choose one of these topologies according to possible special criteria regarding the cost and the network topology.

Another important part of the output is information about routing of channels in the normal and failure states. This output is given in tabular form; Fig. 8 shows a part of the routings table for the example of Figs. 4 and 5 with the capacities shown in Fig. 6. Figure 9 shows the routing of a demand in



■ Figure 5. Logical switching network (demand graph). The demands are given in number of 64 kb/s channels.

the case of diversification for the example of Figs. 4 and 5 and the capacities of Fig. 7.

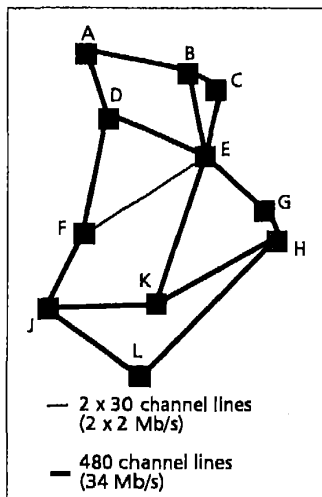
For every operating state (normal, node failure, and link failure) the routing table provides information of routings and reroutings, and may be used as a database for network management systems which are involved in network recovery.

## IMPLEMENTING SURVIVABILITY


The model presented in this article gives the network designer various ways to introduce survivability at the transport level. In this section we discuss and compare these.

The physical network is said to have survivability of  $\alpha$  percent if at least  $\alpha$  percent of each demand can be satisfied in case of a single node or single edge failure. In our model we have two input parameters which are used to introduce survivability: the diversification and reservation parameters. These can be set one at a time or in any combination.

Setting the diversification parameter  $\delta_{uv}$  for the demand  $d_{uv}$  of the logical switching network, we require at most  $100\delta_{uv}$  percent of  $d_{uv}$  to be routed through any node (other than  $u$  and  $v$ ) or any link of the physical network. This implies that we get routings which provide node disjoint paths, each carrying at most  $\delta_{uv}d_{uv}$  channels; therefore, only that many channels of the demand can be lost in a single node or single link failure. That is,  $(1 - \delta_{uv})d_{uv}$  channels survive without any rerouting effort. There are two drawbacks, however. First, setting the diversification parameter to  $\delta_{uv}$  implies that the demand  $d_{uv}$  will be routed through at least  $(1/\delta_{uv})$  node disjoint paths. For example, setting  $\delta_{uv}$  to 0.49, we "request" at least three node disjoint paths through which we route  $d_{uv}$ . Evidently, we cannot achieve 100 percent survivability with this parameter. Moreover, diversification values below 0.34 are undesirable by the network operator, because this would force at least four paths, each carry-



■ Figure 6. Resulting physical network with a planned survivability of 50 percent for a single link or node failure.



■ Figure 7. Calculated topology applying a diversification of 50 percent.


Normal operating state:

Demand (K-L)	
Flow	Path nodes
131	K - H - L
127	K - J - L
Demand (H-L)	
Flow	Path nodes
115	H - L
...	
Demand (F-J)	
Flow	Path nodes
212	F - J

Failing edge (H - L):

Demand (K-L)	
Flow	Path nodes
129	K - E - D - F - J - L
Demand (H-L)	
Flow	Path nodes
58	H - G - E - D - F - J - L
...	
Demand (F-J)	
Flow	Path nodes
93	F - J
13	F - D - A - B - E - K - L

■ Figure 8. Example of the routing information provided by DISCNET.



■ Figure 9. In this example the demand (A-F) is routed via three different paths under the condition that no more than 50 percent of the total demand (108 channels) is carried by one path.

survivability values 0, 25, 50, 66, 75, and 100 percent, where the last value can be achieved only by setting the reservation parameter equal to 1.0. Although 75 percent survivability can be achieved by setting the diversification parameter equal to 0.25, we do not consider this option because, as we mentioned above, this forces too many paths for each demand. For the other values of survivability, the corresponding diversification/reservation values are 1.0/0.0, 0.75/0.25, 0.5/0.5, and 0.34/0.66. The best solution values we get with our network dimensioning tool, DISCNET, are shown in Fig. 10.

In general, there is a trade-off between the ease of network management provided by the first method and the low total installation cost provided by the second. Since the network management costs are not included in the installation costs, it is up to the network operator to decide

ing only a small fraction of the demand. The second drawback is the high cost of the resulting network (Fig. 10).

Using the reservation parameter to introduce survivability, we take advantage of possible redundancy in the network by allowing rerouting in failure situations. Depending on the particular failure, all demands might be rerouted. For a specific demand of  $d_{uv}$  channels, the reservation parameter  $r_{uv}$  guarantees that at least  $r_{uv} d_{uv}$  channels will be still satisfied in a failure state. In our tests we have observed that many more can actually survive. For instance, by maximizing — in a post-processing step — the total satisfied demand, we found that all but a few demands are indeed fully satisfied. However, it should be noted that this is an empirical observation, and in theory one can guarantee only that  $r_{uv} d_{uv}$  channels will survive a failure. The advantage of this method is the low cost of the network compared to that of the previous method (Fig. 10). The obvious disadvantage of this method is the need for rerouting in case of a failure. Indeed, as we have observed in practice, this rerouting may be extensive, making management of the network rather difficult.

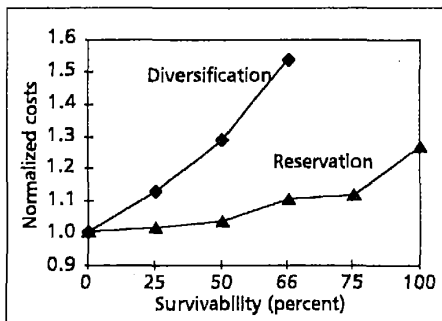
Therefore, one has two methods of introducing survivability to the network, namely, by setting the diversification parameter and by setting the reservation parameter. To compare the costs of the two methods, we make several runs for the example shown in Figs. 4 and 5. We choose as

whether these costs counterbalance the difference in installation costs.


A third method we consider to introduce survivability in the network is a combination of the first two. A minimum survivability is achieved by the diversification parameter setting, with the advantage of easy network management. Additional survivability is introduced by the reservation parameter setting. In case of a failure situation the operator has to decide whether to reconfigure the network or not. This decision depends on various aspects, such as the affected traffic, expected recovery time, and effort required to reconfigure the network.

To compare the cost of implementing the third method to those of the previous ones, we run two additional series of tests, combining diversification and reservation parameters. In the first series we keep a minimum survivability of 25 percent (achieved by setting the diversification parameter to 0.75) and increase survivability by setting the reservation parameter to 0.50, 0.66, 0.75, and 1.0. In the second series we change the minimum survivability value to 50 percent and increase survivability by setting the reservation parameter to 0.66, 0.75, and 1.0. We only consider reservation parameter settings bigger than the minimum survivability, since the diversification parameter setting dominates the other cases.

The best solution values we get



■ Figure 10. Comparison of costs of introducing survivability by setting the diversification or reservation parameter.



■ Figure 11. Comparison of costs for the different methods of introducing survivability.

with DISCNET are shown in Fig. 11. The third curve in Fig. 11 is the reservation curve of Fig. 10 (minimum survivability of 0 percent).

## CONCLUSIONS

Costs, quality, and operational aspects are among the most important issues to be considered in the process of designing a telecommunications network efficiently. The presented network planning tool, DISCNET, provides different low-cost network topologies that guarantee a specified survivability.

We show that the two applicable protection schemes, *diversification* and *reservation*, lead to network topologies that differ in transmission costs and the effort necessary to manage the network. In general, the reservation method leads to networks with lower transmission costs, while the diversification method leads to networks that are easier to manage. The combination of the two protection schemes gives the network operator the opportunity to evaluate the importance of the failure of a network component and to react accordingly.

Taking into account the ongoing liberalization process of the communication environment, efficiently designed telecommunications networks are very important both now and in the future. For this reason planning tools like DISCNET are needed, not only to calculate actual networks, but also to provide insight when used in studies of different network scenarios.

## REFERENCES

- [1] M. Minoux, "Optimum Synthesis of a Network with Non-Simultaneous Multicommodity Flow Requirements," P. Hansen, ed., *Studies on Graphs and Discrete Programming*, North-Holland, 1981, pp. 269-77.
- [2] A. Lissner, R. Sarkissian, and J.-P. Vial, "Optimal Joint Synthesis of Base and Reserve Telecommunications Network," Tech. rep., Univ. Geneva, Oct. 1995.
- [3] D. Bienstock and O. Günlük, "Computational Experience with a Difficult Mixed-Integer Multicommodity Flow Problem," *Math. Programming*, vol. 68, 1995, pp. 213-37.
- [4] T. L. Magnanti, P. Mirchandani, and R. Vachani, "The Convex Hull of Two Core Capacitated Network Design Problems," *Math. Programming*, vol. 60, 1993, pp. 233-50.
- [5] T. L. Magnanti, P. Mirchandani, and R. Vachani, "Modeling and Solving the Two-Facility Capacitated Network Loading Problem," *Ops. Res.*, vol. 43, no. 1, 1995, pp. 142-56.
- [6] G. Dahl and M. Stoer, "MULTISUN — Mathematical Model and Algorithms," Tech. rep. TF-R-46/92, Televerkets Forskningsinstitut, Norway, 1992.
- [7] M. Stoer and G. Dahl, "A Polyhedral Approach to Multicommodity Survivable Network Design," *Numerische Mathematik*, vol. 68, 1994, pp. 149-67.
- [8] G. L. Nemhauser and L. A. Wolsey, *Integer and Combinatorial Optimization*, New York: Wiley, 1988.
- [9] M. Padberg, *Linear Optimization and Extensions*, Berlin: Springer-Verlag, 1995.
- [10] M. Grötschel, C. L. Monma, and M. Stoer, "Computational Results with a Cutting Plane Algorithm for Design Communications Networks with Low-Connectivity Constraints," *Ops. Res.*, vol. 40, no. 2, 1992, pp. 309-30.
- [11] M. Grötschel, C. L. Monma, and M. Stoer, "Facets for Polyhedra Arising in the Design of Communication Networks with Low-Connectivity Constraints," *SIAM J. Optimization*, vol. 2, no. 3, 1992, pp. 474-504.
- [12] M. Stoer, "Design of Survivable Networks," *Lecture Notes in Mathematics*, No. 1531, Berlin: Springer-Verlag, 1992.
- [13] D. Alevras, M. Grötschel, and R. Wessály, "Dimensioning of Survivable Telecommunications Networks," Preprint SC-96-49, Berlin: Konrad-Zuse-Zentrum, 1996.
- [14] M. Irl, "On an Extension of the Maximum-Flow Minimum-Cut Theorem to Multicommodity Flows," *J. Ops. Res. Soc. Japan*, vol. 13, 1971, pp. 129-35.
- [15] K. Onaga and O. Kakusho, "On Feasibility Conditions of Multicommodity Flows in Networks," *Trans. Circuit Theory*, vol. 18, 1971, pp. 425-29.
- [16] S. Martello and P. Toth, "Knapsack Problems," *Algorithms and Computer Implementations*, Chichester, UK: Wiley, 1990.

## BIOGRAPHIES

DIMITRIS ALEVRAS received his diploma in chemical engineering from the National Technical University of Athens, Greece, in 1988, and his M.Sc. and Ph.D. in statistics and operations research from New York University in 1991 and 1994, respectively. In 1994 he joined the Combinatorial Optimization group at the Konrad-Zuse-Zentrum für Informationstechnik Berlin, Germany, where he was involved in research in the area of survivable telecommunications networks design. Since 1996 he has been with the Supply Chain Optimization Practice of IBM Corporation.

MARTIN GRÖTSCHEL received his diploma in mathematics from the University of Bochum (Germany) in 1973 and his Ph.D. from the University of Bonn in 1977. He was a full professor of applied mathematics at the University of Augsburg from 1982 to 1991. He moved to Berlin in 1991 and has been a professor of information technology at the Technische Universität and vice president of Konrad Zuse Zentrum für Informationstechnik Berlin since. His research interests include optimization, operations research, discrete mathematics, and applications thereof. He has been active in several telecommunication applications since 1987, in particular in projects where combinatorial optimization techniques can be employed.

PETER JONAS (jonas@eplus.de) received a diploma and a Ph.D. from the Technical University of Braunschweig, Germany, in 1986 and 1994, respectively. From 1986 to 1992 he was involved in ATM network studies for real-time services like audio and video. In 1993 he joined the fixed network design department of the third German mobile network operator, E-Plus-Mobilfunk GmbH, in Düsseldorf, where he was involved in tool-supported network design. Since 1996 he has been responsible for traffic and network analysis at E-Plus.

ULRICH PAUL received a diploma from the University of Paderborn, Germany, in 1989. From 1990 to 1993 he was involved in the standardization process for a terrestrial flight telephone service (TFTS) and international GSM projects on behalf of DeTeCon GmbH and Deutsche Bundespost Telekom, respectively. In 1993 he joined E-Plus Mobilfunk GmbH in Düsseldorf to work in the fixed network design department to build up the DCS1800 mobile network. In July 1996 he moved to the German private network operator o.tel.o Communications GmbH & Co., where he is now involved in regulatory issues and economics.

ROLAND WESSÁLY received his diploma in computer science from the Technical University of Berlin in 1994. He joined the Konrad-Zuse-Zentrum für Informationstechnik Berlin in 1994. His main research interests are in discrete mathematics, optimization, and operations research. He is currently working on his Ph.D. in the design of survivable telecommunication networks.