

Design of Broadband Virtual Private Networks: Model and Heuristics for the B-WiN

Andreas Bley, Martin Grötschel, and Roland Wessäly

ABSTRACT. We investigate the problem of designing survivable broadband virtual private networks that employ the Open Shortest Path First (OSPF) routing protocol to route the packages. The capacities available for the links of the network are a minimal capacity plus multiples of a unit capacity. Given the directed communication demands between all pairs of nodes, we wish to select the capacities in a such way that even in case of a single node or a single link failure a specified percentage of each demand can be satisfied and the costs for these capacities are minimal. We present a mixed-integer linear programming formulation of this problem and several heuristics for its solution. Furthermore, we report on computational results with real-world data.

1. Introduction

In this article we describe the network design and bandwidth allocation problem we encountered in a cooperation with our partner DFN-Verein e.V. (Registered Association for the Promotion of a German Research Network). DFN acts as the internet provider for German universities and research institutions. The backbone network of DFN, called B-WiN, is operated as a broadband virtual private network (BVPN) on the ATM cross connect network of the Deutsche Telekom AG.

Currently, the network contains ten central service switches. It is possible to rent a link at a certain capacity between each pair of switches. These links are virtual paths in the ATM cross connect network, whose structure is transparent to DFN. The capacities available for rented links are a certain minimum capacity plus multiples of a unit capacity. See [11] for a description of the technical implementation of the B-WiN.

In the network design problem considered here, we must employ the OSPF (Open Shortest Path First, see [7] and [23]) routing protocol to route the traffic demands in the network. Using this routing policy, each package is sent to its destination along a shortest path with respect to some given link weights. In principle, the routing is performed in an embedded IP overlay network consisting

1991 *Mathematics Subject Classification.* Primary 90B12, 90C11; Secondary 90C27, 90C90.

Key words and phrases. Survivable Network Design, Network Capacity Planning, OSPF Routing, Shortest Path Routing, Heuristics.

This work was partially funded by the Bundesministerium für Bildung, Wissenschaft, Forschung und Technologie (BMBF).

of permanent virtual circuits (PVCs), which are paths in the BVPN. However, since DFN currently identifies virtual paths and virtual circuits, we do not distinguish either.

Nowadays, a network provider must not only offer large bandwidths and high transmission rates, but also guarantee a specified quality of service. Hence, it is important for the provider to limit the impact of network failures. In this article we consider single failures of switches or links. Given the directed traffic demands between all pairs of switches, the capacities must be selected in such a way that even in case of a single component failure at least a specified percentage of each demand can be routed in the remaining network.

The design of survivable networks has been studied for various capacity and survivability models. Grötschel, Monma, and Stoer considered uncapacitated network design problems with connectivity constraints in [9, 10, 22]. The networks synthesis problem which includes some variation of capacity, routing, and survivability decisions attracted many researchers. Gomory and Hu initiated the research on non-simultaneous single-commodity flow problems in [8]. Minoux addressed in [18] non-simultaneous multi-commodity flow problems, suggesting a sub-gradient optimization procedure for its solution. More recently, link and path restoration have been considered, for instance, in [2, 15, 19, 25]. The solution methods include combinatorial heuristics, LP-based rounding heuristics, and branch-and-cut algorithms. The survivability concepts of diversification and reservation have been investigated in [1] and [4] for two different capacity models and cutting plane methods have been suggested as solution approach. Multi-layer models including survivability have been considered, for instance, in [13, 17]. Given a network topology and link capacities, Lin and Wang addressed in [16] the problem of finding a set of link weights that provide a feasible OSPF routing and proposed a dual solution approach based upon Lagrangean relaxation.

In this article we present a model that differs in one important aspect from the models studied previously. We integrate survivability, capacity planning and the OSPF routing protocol in one model. Survivability respecting the OSPF routing protocol can be viewed as a variation of path restoration. The failing communication demands are rerouted end-to-end, but the way the routing paths may change in a failure situation is defined by the routing protocol and each demand is routed on a single path. Bifurcated routing is allowed neither in the non-failure nor in any of the failure cases.

The remainder of this article is organized as follows: The next section contains a detailed description of the problem, its input data, and its constraints. A mixed-integer linear programming model for the IP network design and routing problem is developed in Section 3. In Section 4 we briefly present several heuristics to solve this problem. Computational results for real-world data provided by DFN are reported in the last section.

2. The Problem

Formally, the problem can be described as follows: We are given a **supply graph** $G = (V, E)$ consisting of the **nodes** V and the **edges** E . The node set V corresponds to the set of IP router locations. The edge set E corresponds to the set of all potential links, which in our application is the set of all virtual paths (in the underlying ATM-network) that may be rented from Deutsche Telekom.

For every supply edge, the installed capacity must be either zero or the **minimal capacity** $c_{\min} \in \mathbb{Z}_+$ plus a multiple of the **unit capacity** $c_u \in \mathbb{Z}_+$. For no edge may the capacity exceed the **maximal capacity** $c_{\max} \in \mathbb{Z}_+$. The capacities are bidirectional, i.e., the capacity installed on an edge can be used in both directions independently. An edge with non-zero capacity is called a **chosen edge**.

We wish to design a network that is still operational if a single node or a single edge fails. Therefore, we introduce the set of **operating states** S . The considered operating states $s \in S$ are the **normal operating state** ($s = 0$), which is the state with all nodes and edges operational, and the **failure states**, which are the states with a single edge ($s = e \in E$) or a single node ($s = v \in V$) non-operational. Note that in a node failure state $s = v \in V$ all edges incident to v are non-operational, too.

For each ordered pair of nodes $(u, v) \in V \times V$ the value $d_{uv} \in \mathbb{Z}_+$ is the directed **communication demand** from u to v . Clearly, in a node failure state $s = v \in V$ the demands with origin or destination v cannot be satisfied and are therefore not considered in this operating state.

Since network failures are considered to be non-permanent, we do not have to satisfy the total demands in a failure situation. For each operating state $s \in S$, we introduce a **reservation parameter** $r^s \geq 0$ specifying how many percent of each demand must be routable in this operating state. In contrast to telephone networks, the traffic demands in IP networks are extremely bursty and the network contains a higher percentage of spare capacity to accommodate the traffic peaks. Our model allows to oversize the installed capacities in the normal operating state by setting the reservation parameter r^0 to a value larger than one. In this case, even if all demands increase simultaneously by a factor of r^0 , the capacities permit a feasible routing as long as no network component fails.

In principle, the demands are routed in a second network layer, the so called overlay network, which is embedded in the BVPN. The nodes of the overlay network are the nodes of the supply graph. Its edges, which are permanent virtual circuits (PVCs) in the BVPN, correspond to paths in the supply graph. For each PVC a capacity is reserved on its edges. With this second network layer it is easier to manage different traffic types, such as IP, X.25, etc., in the network. Usually, one sets up a separate overlay network for each traffic type and routes the corresponding demands within this overlay network. However, we assume here that there is only one traffic type in the network, namely IP traffic. Furthermore, we restrict ourselves to the case where, for each rented virtual path, there is exactly one direct permanent virtual circuit set up at the full capacity of the virtual path, i.e., the IP overlay network and the BVPN are equivalent. Hence, we can pretend that the demands are routed directly in the supply graph and *no* second network layer exists.

The capacities chosen for the supply edges have to be large enough to allow a feasible routing with respect to the OSPF routing protocol in all operating states. Assuming non-negative **routing weights** for all supply edges, the OSPF protocol implies that each demand is sent from its origin to its destination along a shortest path with respect to these weights. In each operating state only operational nodes and edges are considered in the shortest path computation. In this article we address only static OSPF routing where the link weights do not depend on nor change with the traffic. Dynamic shortest path routing algorithms, which try to adapt to traffic changes, often cause oscillations that lead to significant performance degradation, especially if the network is heavily loaded (see [3]).

Although not required by the OSPF protocol, in our problem we have to guarantee the existence of a path between any pair of operational nodes in each operating state. This implies that the subgraph of G induced by the chosen supply edges must be 2-node-connected.

The routing weights must be chosen in such a way that, for all operating states and all demands, the shortest path from the demand's origin to its destination is unique with respect to these weights. Otherwise, it is not determined which one of the shortest paths will be selected by the implementation of the routing protocol in the network and, therefore, it would be impossible to guarantee that the chosen capacities permit a feasible routing of the demands.

For each demand, the variation of its data package transmission times increases significantly with the number of nodes in the routing path, especially if the network is heavily loaded. For multi-media applications, for example, this might lead to unacceptable differences in the transmission times. In order to avoid too large variations, the number of edges in each routing path is bounded from above by $\ell \in \mathbb{Z}_+$ in the normal operating state.

Finally, let us explain the (complicated) cost structure of the problem considered. Since the edges of the network correspond to virtual paths in the ATM-provider's network, we have network costs which do not depend on the physical lengths of the edges but on the number of chosen edges. Hence, we cannot allow an arbitrary number of edges in a solution. Due to the different tariffs for different numbers of edges, we have a predefined number, say m , of supply edges that must be chosen. All the following cost parameters depend on this number m . Router costs are not considered. In our application each node is a central service switch for its local area and, hence, must have routing equipment installed anyway.

The total cost of installing capacities on the supply edges is defined as follows. For installing the minimal capacity c_{min} on m edges we have the **fixed cost** $\mathbf{K}_{min} \in \mathbb{Z}_+$. Since we cannot influence these costs, they can be omitted in the optimization.

Each capacity unit c_u that is installed in addition to the minimal capacity on a single supply edge causes further costs. The set of feasible capacities above the minimal capacity c_{min} is divided into consecutive capacity intervals (batches). Each interval (except maybe the last interval) contains $\mathbf{b} \in \mathbb{Z}_+$ capacity units and, thus, has a **batch capacity** $c_{batch} := \mathbf{b} \cdot c_u$. The number of these intervals is $\lceil \frac{c_{max} - c_{min}}{c_{batch}} \rceil$. For each interval $i \in I := \{0, \dots, \lceil \frac{c_{max} - c_{min}}{c_{batch}} \rceil - 1\}$, we denote the cost for one capacity unit c_u installed within $[c_{min} + i \cdot c_{batch}, c_{min} + (i + 1) \cdot c_{batch}]$ by $\mathbf{k}_i \in \mathbb{Z}_+$.

For one capacity unit installed *simultaneously on all* m chosen supply edges in interval $i \in I$, the sum of the unit costs of these m capacity units is reduced by the **discount** $\mathbf{K}_i \in \mathbb{Z}_+$. For notational convenience, we denote the minimum of the capacities of the m chosen supply edges as **network base capacity** c_{base} . Note that the cost discount applies exactly to those capacity units that are installed below the network base capacity.

Let us illustrate the capacity and cost structure by means of the example shown in Figure 1. There are six supply edges, five of which are chosen. The minimal capacity c_{min} and the unit capacity c_u are 80,000 and 5,000 ATM-cells/s, respectively. The batch capacity is equal to the minimal capacity, i.e., $c_{batch} = 80,000$ ATM-cells/s and $b = 16$. In this example we have a base capacity c_{base} of 130,000 ATM-cells/s. The cost for installing the minimal capacity of 80,000 ATM-cells/s

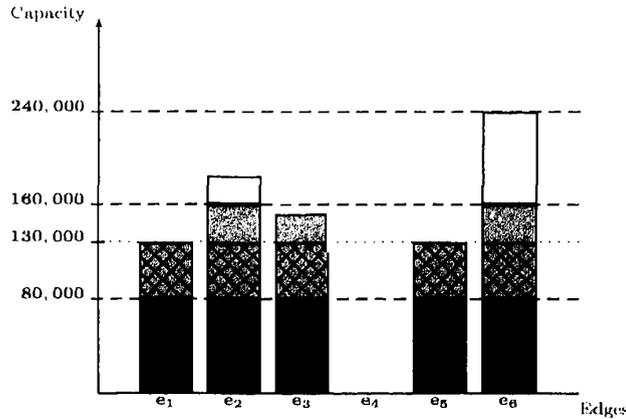


FIGURE 1. Example with six potential links

(dark gray) on all five chosen edges is the fixed cost K_{min} . For every capacity unit installed above 80,000 ATM-cells/s we have costs that depend on the interval in which the unit is installed. Up to 160,000 ATM-cells/s (mid gray) we have costs of k_0 per unit and between 160,000 and 240,000 ATM-cells/s (light gray) we have costs of k_1 per unit. For the $130,000 - 80,000 = 10 \cdot 5,000$ ATM-cells/s installed additionally to the minimal capacity on all five chosen supply edges simultaneously (hatched area) we get a discount of $10 \cdot K_0$.

Note that in the problem considered here the cost of a solution does not depend on the physical lengths of the chosen edges.

The **IP network design problem**, called **IP-ND**, consists of two parts. In the first part, we design the overall BVPN, i.e., we decide which capacities to install on which edges. In the second part, we determine the routing weights and compute the routings for each operating state. The objective is to minimize the total costs of the network, which depend on the installed capacities.

3. The Model

In this section we present a **mixed-integer linear programming model** of the IP-ND network dimensioning and routing problem.

3.1. The Supply Graph. To describe which edges are chosen at which capacities we have so-called global and local variables in our model. The global variables are used to decide the network base capacity, whereas the local variables are employed to decide whether a supply edge will be installed in a solution, and if so, to choose the capacity installed on this particular edge. The **global variables** are

- $x_i \in \{0, \dots, b\}$, for all $i \in I$, determining the number of capacity units installed in interval i as base capacity.

The **local variables**, for every supply edge $e \in E$, are

- $z_e \in \{0, 1\}$, where $z_e = 1$, if we install capacity on supply edge e ,
- $z_{i,e} \in \{0, 1\}$, for all $i \in I$, where $z_{i,e} = 1$ if the capacity of edge e is larger than $c_{min} + i \cdot c_{batch}$, and
- $x_{i,e} \in \{0, \dots, b\}$, for all $i \in I$, determining the number of capacity units installed in interval i on edge e .

To represent feasible choices, the local variables associated with a supply edge $e \in E$ must satisfy

$$\begin{aligned} (1) \quad & z_{i,e} - z_e \leq 0 \quad \forall i \in I, \\ (2) \quad & x_{i,e} - b \cdot z_{i,e} \leq 0 \quad \forall i \in I, \text{ and} \\ (3) \quad & x_{i,e} - b \cdot z_{i+1,e} \geq 0 \quad \forall i \in I. \end{aligned}$$

Due to inequality (1) the variable $z_{i,e}$ can only be positive if $z_e = 1$. Suppose we wish to choose edge e in a solution, i.e., we have $z_e = 1$. In this case it follows from (2) and (3) that the variables $x_{i,e}$ can only be positive if $z_{i,e} = 1$ and that $z_{i+1,e} = 1$ implies $x_{i,e} = b$. Hence, if capacity units are installed in interval $i \in I$, then $z_{j,e} = 1$ and $x_{j,e} = b$ for all $j < i$, i.e., all b capacity units are installed in the capacity intervals “below” i , and $z_{i,e} = 1$. This determines a proper setting of the local variables and the edge’s capacity. If we do not choose edge e , i.e., $z_e = 0$, then the inequalities (1), (2), and (3) imply $z_{i,e} = x_{i,e} = 0$ for all $i \in I$.

Note that, if the capacity of edge e equals $c_{min} + i \cdot c_{batch}$ for some $i \in I$, then the variable $z_{i,e}$ may be zero as well as one.

With the following constraints we guarantee that the capacity of each chosen supply edge exceeds the base capacity:

$$(4) \quad x_i \leq x_{i,e} + b(1 - z_e) \quad \forall i \in I, e \in E.$$

To understand (4) consider the following. If supply edge e is chosen in the network, then $z_e = 1$. In this case inequality (4) reduces to $x_i \leq x_{i,e}$, which means that the number of capacity units installed in interval i on edge e is at least the number of units installed for the network base capacity. In the other case $z_e = 0$ and (4) is satisfied anyway.

For notational convenience, we introduce auxiliary variables

$$(5) \quad y_e := c_{min} \cdot z_e + c_u \sum_{i \in I} x_{i,e} \quad \forall e \in E,$$

representing the capacities installed on the supply edges.

It is easy to see that, if $K_i > 0$ for all $i \in I$, in any optimal solution of our mixed-integer linear programming model the term $\sum_{i \in I} x_i$ is the network base capacity. The upper bounds for these capacities are enforced by

$$(6) \quad y_e \leq c_{max} \quad \forall e \in E.$$

To ensure that the number of chosen supply edges is exactly m we have the equality

$$(7) \quad \sum_{e \in E} z_e = m.$$

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

$$(8) \quad \min \quad K_{min} + \sum_{e \in E} \sum_{i \in I} k_i x_{i,e} - \sum_{i \in I} K_i x_i.$$

The term $K_{min} + \sum_{e \in E} \sum_{i \in I} k_i x_{i,e}$ describes the total cost of the installed capacity units and $\sum_{i \in I} K_i x_i$ is the discount for those capacity units that are installed for the network base capacity.

3.2. The OSPF-Routing. In the following we present a model of the OSPF routing protocol, which is used to route the IP traffic in the network.

Let $G^s = (V^s, E^s)$ be the supply graph in operating state $s \in S$, i.e., the subgraph of G consisting of all nodes and edges that are still operational. Since the IP traffic is directed, we associate with each G^s the directed graph $D^s = (V^s, A^s)$ with

$$A^s := \bigcup_{e \in E^s} \{\bar{e}^s, \overleftarrow{e}^s\},$$

where \bar{e}^s and \overleftarrow{e}^s denote the two orientations of edge e . By $D = (V, A) = (V^0, A^0)$ we denote the digraph associated with G .

Using the OSPF routing protocol, each package is routed to its destination on a shortest path with respect to a common edge weight function. All packages emanating from a node with the same destination must use the same route, bifurcation is not allowed. In our model we use the **path variables**

- $t_{uv,a}^s \in \{0, 1\}$, for all $s \in S$, $u, v \in V^s$, and $a \in A^s$, which are chosen to be one if arc a is contained in the routing path from u to v in operating state s .

To guarantee that, for every operating state $s \in S$ and every pair $u, v \in V^s$ of nodes, the arcs corresponding to the non-zero path variables form a path we introduce the equalities and inequalities

$$(9) \quad \sum_{a \in \delta_{D^s}^+(u)} t_{uv,a}^s = \sum_{a \in \delta_{D^s}^-(v)} t_{uv,a}^s = 1,$$

$$(10) \quad \sum_{a \in \delta_{D^s}^+(w)} t_{uv,a}^s - \sum_{a \in \delta_{D^s}^-(w)} t_{uv,a}^s = 0 \quad \forall w \in V^s \setminus \{u, v\}, \text{ and}$$

$$(11) \quad \sum_{a \in A^s(W)} t_{uv,a}^s \leq |W| - 1 \quad \forall \emptyset \neq W \subseteq V^s,$$

where $A^s(W)$ denotes the set of arcs in the subgraph of D^s induced by W and $\delta_D^+(u)$ and $\delta_D^-(u)$ denote the sets of arcs in D with tail and head u , respectively. Equality (9) ensures that there is exactly one arc leaving the source node u and one arc entering the destination node v . For every other node we ensure that the number of incoming arcs equals the number of leaving arcs by (10), while (11) eliminates additional cycles.

To guarantee that no routing path contains more than ℓ arcs in the normal operating state we need the inequalities

$$(12) \quad \sum_{a \in A^0} t_{uv,a}^0 \leq \ell \quad \forall u, v \in V.$$

For each operating state $s \in S$, the installed capacities must permit a feasible routing of the specified reservation percentage r^s of each demand. This condition yields the capacity constraints

$$(13) \quad r^s \sum_{u,v \in V^s} t_{uv,a}^s d_{uv} \leq y_e \quad \forall s \in S, e \in E^s, a \in \{\bar{e}^s, \overleftarrow{e}^s\}.$$

Up to here, we have modeled a survivable single path routing scheme, i.e., a scheme where in each operating state there is exactly one routing path between any pair of nodes. To model the OSPF routing protocol completely, we also have to ensure that all routing paths are shortest paths with respect to a common weight function $w : E \rightarrow \mathbb{R}_+$. We introduce the following variables:

- $w_e \in \mathbb{R}_+$, for each $e \in E$, which is the **routing weight** of edge e , and
- $\pi_{v,u}^s \in \mathbb{R}_+$, for each $s \in S$ and each pair $u, v \in V^s$, the **potential variables** denoting a feasible potential of node v with respect to the edge weights w and the root node u in D^s .

The node potentials are used to decide which arcs are on shortest paths and can be used in routing paths and which are not. They correspond to the dual variables in the embedded shortest path problems (see [6, 12]).

The feasibility of the potentials is guaranteed by the following metric inequalities:

$$(14) \quad \pi_{u,u}^s = 0 \quad \forall u \in V^s,$$

$$(15) \quad \pi_{w,u}^s - \pi_{v,u}^s + w_e \geq 0 \quad \forall u, v, w \in V^s, vw = e \in E^s.$$

It follows from (14) and (15) that in each solution of this mixed-integer program the variable $\pi_{v,u}^s$ is at most the distance from v to u in G^s with respect to w . Note that the perturbation technique that will be introduced below to ensure uniqueness of all shortest paths implies that all routing weight variables are strictly positive.

If, for some feasible potentials, inequality (15) holds with equality for all edges of a path with destination u , then this path is a shortest path from its origin to u in G^s (see [6, 5]). Furthermore, there exist feasible potentials, such that edge e is contained in a shortest path with destination u in G^s if and only if inequality (15) is tight. Hence, we can ensure that all routing paths are shortest paths from their origins to their destinations, by introducing for each operating state $s \in S$ and each pair of operational nodes $u_1, u_2 \in V^s$ the inequalities

$$(16) \quad \pi_{w,u_2}^s - \pi_{v,u_2}^s + w_e + M(t_{u_1 u_2, a}^s - 1) \leq 0 \quad \forall (v, w) = a \in A^s, e = vw,$$

where M is a large constant. It is sufficient to choose M to be two times the maximum of the routing weights.

Finally, we have to ensure that all routing paths are *unique* shortest paths with respect to w . We apply a perturbation technique that guarantees the uniqueness of the lengths of *all* paths in G . Let $\sigma : E \rightarrow \{1, \dots, |E|\}$ be an arbitrary permutation of the edges of G . We define the **perturbation** of the routing weights as

$$(17) \quad \epsilon_e : E \rightarrow \mathbb{R}_+, \quad \epsilon_e := 2^{\sigma(e) - |E| - 1}.$$

It is easy to see that $\sum_{e \in E} \epsilon_e < 1$. Furthermore, for all subsets $E_1, E_2 \subseteq E$ we have $\sum_{e \in E_1} \epsilon_e \neq \sum_{e \in E_2} \epsilon_e$ if and only if $E_1 \neq E_2$.

Let $w'_e : E \rightarrow \mathbb{Z}_+$ be *integer* weights on the edges of G . Then with

$$(18) \quad w : E \rightarrow \mathbb{R}_+, \quad w_e := w'_e + \epsilon_e$$

different paths have different lengths and, thus, every shortest path is unique. This follows directly from $\sum_{e \in P} \epsilon_e < 1$ for all paths $P \subseteq E$, $\sum_{e \in P_1} \epsilon_e \neq \sum_{e \in P_2} \epsilon_e$ for all distinct paths $P_1, P_2 \subseteq E$, and the integrality of w' . Note that the integer weights w' are the variables in this mixed-integer programming model. The perturbed weights w are “only” artificial variables.

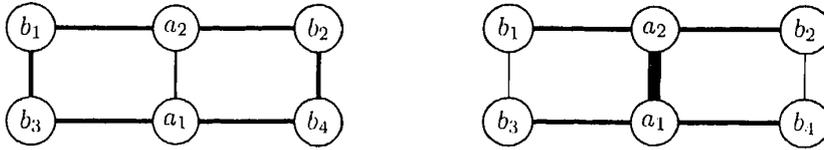


FIGURE 2

We wish to stress at this point the indispensability of all shortest paths being unique. Only if this condition is satisfied, the routing paths chosen by the router in practice are guaranteed to be those computed in the planning. Otherwise, the capacities are potentially insufficient for the paths chosen in practice. As an example, consider the network in Figure 2 with communication demands $d_{a_1 a_2} = d_{a_2 a_1} = 1$, $d_{b_j a_i} = d_{a_i b_j} = 1$ for $i = 1, 2$ and $j = 1, \dots, 4$, and $d_{b_i b_j} = 0$ for $i, j = 1, \dots, 4$, that is a_1 and a_2 have a demand of one from and to any other node. Suppose in a solution we have a capacity $\eta_e = 2$ for all supply edges e . It is easy to verify that these capacities are feasible with respect to the routing weights $w_e = 1 + 2^{\sigma(e)-8}$ for any permutation $\sigma: E \rightarrow \{1, \dots, 7\}$ of the supply edges with $\sigma(a_1 a_2) = 7$. For routing weights $\hat{w}_e = 1$, for all supply edges e , all shortest paths with respect to w remain shortest paths with respect to \hat{w} in this example. However, these paths are no longer *unique* shortest paths, and, even worse, for all non-zero demands there is another \hat{w} -shortest routing path containing supply edge $a_1 a_2$. Thus, if the router chooses these latter paths the required capacity on edge $a_1 a_2$ would be three.

Given the perturbed routing weights for the edges of the supply graph, we can easily determine whether, for these weights, there exists a feasible solution satisfying all side constraints. From inequalities (9)–(11), (14), and (15) we obtain the routing paths and with (13) a lower bound for the capacity of each supply edge. We have to validate three constraints: The routing paths in the normal operating state must not contain more than ℓ edges (12), the number of chosen supply edges must be exactly m (7), and, finally, the lower bound for the capacity of each supply edge may not exceed c_{max} (6). If these constraints are satisfied there exists a solution with these routing weights. The capacities of a cheapest such solution can be easily computed from the routing paths.

4. Algorithmic Approach

In this section we describe several starting and improvement heuristics that utilize the observation made at the end of the previous section. Given (perturbed) routing weights for the supply edges, it is easy to decide whether these weights, together with the induced routing paths and capacities, define a feasible solution. In the starting heuristics we try to assign such weights to the supply edges from scratch, while in the improvement heuristics we iteratively exchange or modify these weights to reduce the induced cost.

4.1. Starting Heuristics. We employ a two step approach to compute an initial feasible solution. First we choose exactly m supply edges that induce a 2-node-connected subgraph containing all nodes of the supply graph. Thereafter, in the second step, we assign weights to the supply edges in such a way, that

only these m edges are used to route the demands. If the induced routing paths and capacities satisfy all constraints we have found a feasible solution. Otherwise, the starting heuristic fails. To compute an initial topology, we implemented three methods, one randomized and two deterministic ones.

Random topology: In the first method we iteratively select m supply edges at random until we obtain a 2-node-connected subgraph containing all nodes of G . To test for biconnectivity we use a simple depth-first-search based algorithm due to Tarjan [24]. Applying the random topology method several times, we compute in very short running time many different topologies to continue with.

Tour based topology: The idea, here, is to compute heuristically a Hamiltonian cycle with a double tree heuristic or with Christofides' heuristic (see [20]) and to add further edges until the induced subgraph contains m edges. Since we start with a Hamiltonian cycle, the final subgraph is 2-node-connected and contains all supply nodes. To increase the probability that an edge $e \in E$ is chosen if there is a high demand between its end-nodes u and v we define artificial edge costs $c_e := 1/\max(d_{uv} + d_{vu}, 1)$ for the Hamiltonian cycle computation. In the second step we add the cheapest $m - |V|$ remaining edges with respect to these costs.

Delete heuristic topology: In the third method to compute an initial topology we start with the entire supply graph and iteratively delete edges until there are exactly m edges left. In every iteration we try to delete one of the remaining edges. If, after deletion, the edges left over do not induce a 2-node-connected subgraph containing all nodes, we put this edge back and label it as un-removable. Whenever we cannot remove any further edge we start a backtracking procedure that re-inserts the edge deleted last and reverts all labels set after its deletion.

The final topology depends on the order the edges are considered for deletion. In the beginning and after each successful deletion we (re-)compute the capacities that are induced by (non-perturbed) unit routing weights on the remaining edges (and "infinite" weights on the deleted edges). In each iteration we try to delete the edge with the smallest of these capacities.

Given a feasible network topology, we initially assign a perturbation of the unit weights to its edges. This approach has the advantage that every demand is always routed on a shortest path with respect to the number of edges, a property, which is appreciated by network operators.

With the perturbation technique described in Section 3, such routing weights only depend on the permutation of the supply edges (see (17) and (18)). Given a set $F \subseteq E$ of m supply edges that induce a 2-node-connected subgraph of G containing all nodes, we first choose a permutation $\sigma : F \rightarrow \{1, \dots, m\}$ of these edges. Then we extend this permutation canonically to a permutation $\sigma' : E \rightarrow \{1, \dots, |E|\}$ of all supply edges, with $\sigma(e) = \sigma'(e)$ for all $e \in F$, and set the routing weights as follows:

$$w_e := \begin{cases} 1 + 2^{\sigma'(e)-|E|-1} & e \in F, \\ m + 2^{\sigma'(e)-|E|-1} & \text{otherwise.} \end{cases}$$

Since $w_{e'} \geq m + 2^{m-|E|} > \sum_{e \in F} w_e$ for all $e' \in E \setminus F$ and (V, F) is a 2-node-connected spanning subgraph of G , no edge in $E \setminus F$ is contained in any routing

path in any operating state. Hence, the subgraph of G induced by the routing paths with respect to these weights is indeed (V, F) . Although in practice only the edges in F are considered in the routing path computations, we set the routing weights for the other edges here according to the model we developed in Section 3.

In our implementation we use the following three methods to choose a permutation of the edges in F and to define the routing weights.

Random permutation: In the first method we choose a permutation of the edges in F at random. If the induced capacities do not exceed the capacity bound c_{max} , we have found a feasible solution. We repeat this method several times and choose the cheapest solution found.

Demand based permutation: To reduce the probability that the capacities induced by the initial routing weights are larger than c_{max} , we order the edges in F in increasing order of the demands between their end-nodes. Using this permutation, edges with higher demands between their end-nodes get larger routing weights than those with smaller demands. Hence, if there is more than one path with the same minimal number of edges between the end-nodes of a demand, we select the path, for whose edges the largest of these demands is minimal.

Capacity based permutation: Similar to the previous perturbation method, we again try to balance the flow on the chosen edges. First, we choose for each operating state and for each demand one of the shortest paths with respect to the number of edges in (V, F) . Using these paths to route the demands, we obtain capacities, by which we order the edges in F increasingly.

4.2. Improvement Heuristics. In this subsection we present several improvement heuristics devised to solve the IP-ND problem, which are based on local search techniques. Given a feasible initial solution, we iteratively exchange or modify the routing weights to reduce the costs caused by the induced capacities.

In principle, a local search algorithm is given by a neighborhood and a rating function. The neighborhood function assigns to each solution a set of solutions, that, usually, can be obtained from this solution by a simple modification. The rating function, which assigns some (maybe artificial) costs to each solution, is used as a measure of quality of the solutions. Given a neighborhood and a rating function, a generic local search algorithm works as follows: Starting with the initial solution, in each iteration the algorithm scans the neighborhood of the current solution for a neighboring solution with the best rating, which becomes the current solution in the next iteration. This process is repeated until there is no solution with a better rating than the current solution in the neighborhood.

To solve the IP-ND problem, we implemented two classes of neighborhood functions and three different rating functions.

Switching neighborhoods: Given a solution, we define its k -switching neighborhood as the set of all feasible solutions that can be obtained by exchanging the routing weights of at most k supply edges. Note that exchanging the weights assigned to chosen and non-chosen edges may change the topology of the induced solution. Since a small number of changes to the routing weights causes only few routing paths to change and these changes are easy to compute (see [21]), the switching neighborhoods can be implemented efficiently.

If the routing weights of the initial solution are perturbed unit weights, which implies that all routing paths are minimal with respect to the number of edges in the current topology, so are the weights of the solutions considered by these neighborhoods.

Modification neighborhoods: In contrast to the switching neighborhood, we modify, here, the edge weights depending on the routing paths and the capacities of the current solution.

The l, k -modification neighborhood of a given solution is defined as follows: In the first step, we compute l assignments of routing weights by modifying the weights of the current solutions. Thereafter, in the second step, for each of these l assignments we consider all weight exchanges defined by the k -switching neighborhood. The l, k -modification neighborhood consists of all feasible solutions induced by the weights obtained this way. Practical experiments revealed that the performance of our algorithms improved significantly when the weight modification methods were combined with the weight switching neighborhoods.

In our implementation we use three different kinds of weight modification methods:

Load: Reduces the routing weight on edges with small load (average flow over all operating states / capacity) to use free capacities.

Capacity: Reduces the routing weight on edges with small capacity to increase the usage of low-capacity edges and decrease the usage of high-capacity edges.

Violation: Increases the routing weight on edges whose current capacity is larger than the capacity in a given reference network.

In contrast to the switching neighborhoods, it may happen that in the solutions computed with these neighborhoods the routing paths are not shortest paths with respect to the number of edges if we have perturbed unit routing weights in the initial solution.

To decide which of the solutions in the current neighborhood is chosen for the next iteration in the local search heuristic, the following three rating functions have been implemented:

Cost: $Cost(solution)$, i.e., the solutions are rated by their cost.

Capacity: $\sum_{e \in E} y_e$, i.e., the solutions are rated by their total capacity.

Violation: $\sum_{e \in E} \max\{y_e - c_e, 0\}$ for given capacities $c_e \in \mathbb{R}_+$, i.e., the solutions are rated by their total violation of the capacities of a given reference network.

The first two rating functions can be applied to reduce the cost of the solution—they evaluate the solutions directly by their cost or indirectly via the total capacity. The last rating function, together with a violation based modification neighborhood, is used in routing heuristics, which try to find routing weights that allow a feasible routing in a given network.

5. Computational results

In this section we report on first computational experiments with real-world data provided by our partner DFN. The current backbone network of DFN contains ten central service switches defining the node set of the complete supply graph. We wish to select either 12 or 13 of its 45 supply edges as cheap as possible for different parameter settings. The directed demands between the central service switches are

Topology	Weights	$m = 12$		$m = 13$	
		Cost	Time (sec)	Cost	Time
Random (50x)	Random	14.02	09.39	13.64	05.47
	Demand	14.36	08.57	12.70	05.50
	Capacity	13.31	12.54	12.11	09.24
Tour	Random	16.58	00.34	13.50	00.32
	Demand	14.00	00.34	14.35	00.33
	Capacity	13.47	00.42	13.27	00.41
Delete	Random	12.47	02.96	11.69	02.84
	Demand	12.50	02.98	11.44	02.86
	Capacity	12.50	03.01	11.55	02.88

TABLE 1. Starting heuristics (month 3, $r^s = 1.0 \forall s \in S \setminus \{0\}$, $r^0 = 2.0$), times in seconds, best three solutions in bold face.

peak demands obtained from IP accounting over three consecutive months. Over these months, the demands' ranges are [110, 56, 000], [160, 59, 000], and [170, 64, 000] ATM-cells/s, respectively. The capacities available for the supply edges are 80, 000 ATM-cells/s plus a multiple of the unit capacity 5, 000 ATM-cells/s, the upper capacity bound is 350, 000 ATM-cells/s. In the reported computations the length restriction for the routing paths in the normal operating state is set to $\ell = 4$.

The algorithms described in the previous section have been implemented in C++ using the library LEDA [14]. The computations were performed on a Sun Ultra Enterprise 3000 (Ultra-SPARC processor at 168 MHz), the peak memory usage was 15 MB.

In Table 1 we compare the solution values and running times for the starting heuristics applied to the last demand set (month). We report the results only for one setting of the reservation parameters. In the normal operating state we oversize the network capacities by a factor of two ($r^0 = 2.0$) and in the failure states we wish to route at least 100 percent of each demand ($r^s = 1.0$ for all $s \in S \setminus \{0\}$). The experiments revealed that the delete method to compute an initial topology clearly outperforms the others, no matter which routing weight assignment we employ in the second phase of the starting heuristics. This result was obtained for all other tested settings of the reservation parameters and all other demand sets, too.

We run several local search heuristics to improve the best solution found by each of the starting heuristics in order to find the best combination. Results for different neighborhoods and demand sets are shown in Table 2. In contrast to the results obtained for the starting heuristics, no combination clearly outperformed the others. As one might expect, it is more difficult to obtain good solutions when starting with a random topology, probably because these do not take advantage of the demand structure.

The running times of our heuristics are fairly good, considering the size of the problems. The starting heuristics compute feasible solutions within a couple of seconds. Keeping in mind that the network is re-designed every 2-3 months only, the running times reported for the combinations of starting and improvement heuristics are also satisfactory, even for the larger neighborhoods.

Heuristic	Month 1		Month 2		Month 3	
	Cost	Time	Cost	Time	Cost	Time
Random (50×)						
+ 2-switch	10.69	32	11.22	48	11.69	42
+ 2,10-mod. (load)	9.95	13:29	11.22	22:38	11.07	20:41
+ 2,10-mod. (cap)	9.96	10:22	11.22	21:43	10.88	13:28
+ 3-switch	10.20	21:50	11.00	12:08	11.57	17:09
+ 3,10-mod. (load)	10.20	2:16:25	10.35	6:35:27	11.05	6:24:04
+ 3,10-mod. (cap)	9.88	3:54:17	10.57	2:12:34	11.18	3:55:02
Tour						
+ 2-switch	10.47	57	10.74	42	11.67	42
+ 2,10-mod. (load)	9.85	23:19	10.74	13:02	11.08	19:36
+ 2,10-mod. (cap)	10.47	5:22	10.71	6:42	10.88	13:28
+ 3-switch	9.95	19:28	10.64	13:45	10.84	21:40
+ 3,10-mod. (load)	9.90	2:10:44	10.57	3:48:52	10.77	2:58:27
+ 3,10-mod. (cap)	9.95	4:48:12	10.64	2:07:11	10.61	6:38:29
Delete						
+ 2-switch	10.10	37	11.00	47	11.37	48
+ 2,10-mod. (load)	9.78	22:36	10.35	25:58	10.86	13:37
+ 2,10-mod. (cap)	9.90	18:28	10.81	06:03	11.01	09:32
+ 3-switch	10.15	14:03	10.56	13:15	11.37	12:30
+ 3,10-mod. (load)	9.88	4:49:08	10.35	4:33:49	10.83	6:21:17
+ 3,10-mod. (cap)	9.95	3:24:09	10.56	1:59:05	10.73	5:53:42

TABLE 2. Starting + improvement heuristics ($m = 13$, $r^0 = 2.0$, $r^s = 1.0 \forall s \in S \setminus \{0\}$), times in h:min:sec, best three solutions in bold face.

r^0	r^s	Best solution costs	
		pert. unit weights	arb. weights
2.0	1.0	11.53	11.31
2.0	0.8	10.94	10.79
2.0	0.0	10.33	9.72
1.0	1.0	10.74	10.66
1.0	0.8	9.17	9.17
1.0	0.0	8.14	8.07

TABLE 3. Impact of the reservation parameter setting (month 3, $m = 12$).

Finally, to evaluate the impact of the reservation parameter on the best solution value we run further test series. We see from Table 3 that there is a clear trade-off between the cost and the quality of the resulting solution. Whether we allow arbitrary or only perturbed unit routing weights does not make a considerable difference. It is up to the network designer to compare the different scenarios and to choose a solution that fits additional design requirements best.

6. Conclusions

In this article we introduced a mixed-integer linear programming formulation for a survivable capacitated network design problem. To our knowledge this is the first time that survivability, capacity planning, and the OSPF routing protocol have been integrated and solved within one model.

We have shown that our heuristics produce solutions that satisfy all constraints in reasonable running times. However, in order to evaluate the quality of the solutions, we still need a method to compute good lower bounds to the optimal solution value.

From a practical point of view, there are at least two interesting extensions to our model. First, a virtual circuit layer that does not necessarily identify virtual paths and virtual circuits and, second, the incorporation of different traffic types or IP service classes in this model.

7. Acknowledgments

We would like to thank DFN and BMBF for funding this project and our colleagues at DFN, in particular Dr. G. Hoffmann, K. Ullmann, and S. Schweizer, for fruitful discussions and excellent cooperation.

References

1. D. Alevras, M. Grötschel, and R. Wesäly, *Cost-efficient network synthesis from leased lines*, Annals of Operations Research **76** (1998), 1-20.
2. A. Balakrishnan, T.L. Magnanti, J. Sokol, and Y. Wang, *Modeling and solving the single facility line restoration problem*, Tech. Report OR 327-98, MIT, Operations Research Center, April 1998.
3. J. Crowcroft and Z. Wang, *Analysis of shortest-path routing algorithms in a dynamic network environment*, ACM SIGCOM Computer Communication Review **22** (1992), no. 2, 63-71.
4. G. Dahl and M. Stoer, *A cutting plane algorithm for multicommodity survivable network design problems*, INFORMS Journal on Computing **10** (1998), no. 1, 1-11.
5. M.M. Deza and M. Laurent, *Geometry of cuts and metrics*, Springer, 1997.
6. András Frank, *Connectivity and network flows*, Handbooks of Combinatorics (R. L. Graham, M. Grötschel, and L. Lovász, eds.), Elsevier-Science B.V., Amsterdam, 1995, pp. 111 - 177.
7. M. Gerla, *Protocols and techniques for data communication networks*, ch. 4, Routing and Flow Control, Prentice-Hall, 1981.
8. R.E. Gomory and T.C. Hu, *Synthesis of a communication network*, Journal SIAM **12** (1964), no. 2, 348-369.
9. M. Grötschel and C.L. Monma, *Integer polyhedra associated with certain network design problems with connectivity constraints*, SIAM Journal on Discrete Mathematics **3** (1990), no. 4, 502-523.
10. M. Grötschel, C.L. Monma, and M. Stoer, *Handbooks in Operations Research and Management Science, Network Models*, ch. 10, Design of Survivable Networks, North-Holland, 1995.
11. G. Hoffmann, *B-WiN - The ATM-based high-speed network for the DFN community*, Computer networks and ISDN Systems (1996), no. 28, 1953-1960.
12. U. Huckenbeck, *Extremal paths in graphs*, Akademie Verlag, 1997.
13. K.R. Krishnan, R.D. Doverspike, and C.D. Pack, *Improved survivability with multi-layer dynamic routing*, IEEE Communications Magazine (1995), 62-68.
14. *LEDA—Library of Efficient Data types and Algorithms*, developed at the Max-Planck-Institut für Informatik Saarbrücken, <http://www.mpi-sb.mpg.de/LEDA/leda.html>.
15. K. Lee, K. Park, and S. Park, *Spare channel assignment for DCS mesh-restorable networks*, Proceedings of the 3rd International Conference on Telecommunication Systems: Modeling and Analysis, 1995, pp. 296-307.
16. F.Y.S. Lin and J.L. Wang, *Minimax open shortest path first routing algorithms in networks supporting the SMDS service*, Tech. report, Bell Communications Research, 1993.

17. D. Medhi, *A unified approach to network survivability for teletraffic networks: Models, algorithms and analysis*, IEEE Transactions on Communications **42** (1994), no. 2/3/4, 534-548.
18. M. Minoux, *Optimum synthesis of a network with non-simultaneous multicommodity flow requirements*, Studies on Graphs and Discrete Programming (P. Hansen, ed.), North-Holland Publishing Company, 1981, pp. 269-277.
19. F. Poppe and P. Demeester, *The design of SDH mesh-restorable networks: Problem formulation and algorithm*, Proc. of the 5th International Conference on Telecommunication Systems, Modeling and Design, 1997, pp. 88-97.
20. G. Reinelt, *The traveling salesman*, Springer, 1994.
21. D.R. Shier and C. Witzgall, *Arc tolerances in shortest path and network flow problems*, Networks **10** (1998).
22. M. Stoer, *Design of survivable networks*, Lecture Notes in Mathematics, vol. 1531, Springer, 1992.
23. A.S. Tanenbaum, *Computer networks*, 3 ed., Prentice-Hall, 1996.
24. R.E. Tarjan, *Depth-first search and linear graph algorithms*, SIAM Journal on Computing **1** (1972), 146-160.
25. T.-H. Wu, *Fiber network service survivability*, Telecommunications Library, Artech House, 1992, ISBN: 0-89006-469-5.

KONRAD-ZUSE-ZENTRUM FÜR INFORMATIONSTECHNIK BERLIN (ZIB), TAKUSTR. 7, D 14195
BERLIN, GERMANY
E-mail address: {bley,groetschel,wessaely}@zib.de