

Large-scale hierarchical networks: How to compute an optimal architecture?

Andreas Bley
Thorsten Koch
Roland Wessály

January, 2004

Abstract

In this article, we present a mathematical model and an algorithm to support one of the central strategic planning decisions of network operators: How to organize a large number of locations into an hierarchy of network levels? We propose a mixed-integer program and a Lagrangian relaxation based algorithm to model and solve this planning task. As one big advantage of this approach, not only solutions but also worst-case quality gurarantees can be provided. We present a solution for a G-WiN planning instance of DFN with 759 locations which has been computed in less than 30 minutes and which is (provably) less than 0.5 percent away from optimality.

Contents

1	Introduction	2
2	Planning task	3
2.1	Locations, hardware and capacities	3
2.2	Routing and survivability	4
2.3	Optimization target	4
3	Mathematical model	4
3.1	Topology, hardware and link capacities	5
3.2	Routing	6
3.3	Cost minimization	6
4	Algorithmic Approach	6
5	Computational results	8
5.1	Planning scenario	8
5.2	The results	10
6	Conclusions	11

1 Introduction

The structure of communication networks is influenced by many technological and organizational boundary conditions. Besides cost, geographical issues must be taken into account as well as technical and functional aspects. Larger networks are partitioned in a hierarchical way into backbone and access networks since these are better suited for planning and operation of the network as a whole; see Figure 1 as an example of a three-level hierarchy.

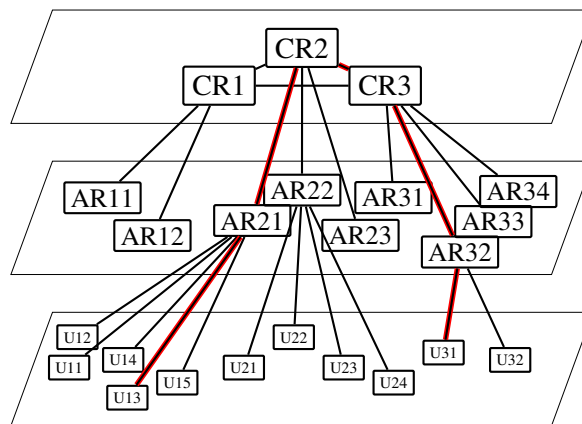


Figure 1: Network with a three-level hierarchy. Core routers (CR) are defined in the backbone, access routers (AR) in the first level, and users (U) in the second level.

As the central strategic planning task, operators of large-scale networks (with several hundreds of locations) must define an appropriate hierarchy, which structures the complete network into manageable and cost-effective sub-networks. A number of questions must be answered in the process of defining this hierarchy:

- How to optimize the trade-off between connection and equipment cost?
- What is the optimal number of backbone locations?
- Which are the backbone locations?
- How to dimension the equipment?
- Which access locations should be attached to which backbone location?

The answers to these question heavily influence all subsequent operational planning issues as well as the cost of the overall network. This type of planning task has to be solved by any operator of a larger communication network, independent of the technology (e.g., IP/OSPF, MPLS, ATM, SDH, WDM). We have been working on similar planning tasks with various project partners, e.g., Telekom Austria as a fixed-network operator, E-Plus Mobilfunk as a mobile-phone network operator, and the DFN-Verein as the provider of the largest German IP-network [5, 6, 10]. Recently, such planning issues arise as a hot

planning topic in UMTS networks, where, such decisions have to be taken for new equipment such as Media Gateways (MGWs) or Remote Network Controllers (RNCs).

With this article we aim at indicating that today's mathematical optimization technology provides automated, intelligent and efficient algorithms to support strategic planning decisions of such complexity. In Section 2, the planning background is described. How to model this in terms of a mathematical model and how to solve the resulting mixed-integer programs by combining a Lagrangian Relaxation with a branch-and-cut algorithm follows in Sections 3 and 4. With Section 5, we demonstrate the power of this algorithm by showing results for a real-world planning instance of DFN with more than 700 locations.

2 Planning task

The DFN-Verein operates the "German Research Network", which connects universities and research institutes from all over Germany and serves as a platform to develop and test new internet applications. The backbone of this network, the so-called "Gigabit-Wissenschaftsnetz" (G-WiN), has been launched in 2000. It has substituted the "Breitband-Wissenschaftsnetz" (B-WiN) since this network was no longer expandable to carry the continuously increasing traffic volume. The G-WiN connects 750 locations and has initially been dimensioned to carry a monthly traffic of approximately 220 Terabyte.

The DFN-Verein aimed at a high-capacity backbone with a relatively small number of locations. All other locations needed to be attached to this backbone using one or more access networks levels. Furthermore, the central role of the backbone enforced a design able to survive at least single link or single node failures. In contrast, for all access networks it was sufficient to have a star- or tree-like structure.

The general planning problem is defined as follows: (i) decide for each location whether it becomes backbone location, (ii) connect all access locations to the backbone, (iii) decide about the topology, the hardware configuration and the capacities of the backbone network and (iv) identify a survivable routing of the communication demands which respects the link capacities. The objective is to minimize the sum of connection and infrastructure cost.

2.1 Locations, hardware and capacities

Starting point is the set of locations, comprising all user locations as sources or destinations of communication traffic. For each location it is specified to which hierarchy levels it might belong to: a location might only be used in the backbone, for instance, if there is sufficient space for the hardware and if the necessary maintenance and service personnel is available. It also might be that some location has to be in the backbone since it serves as gateway to external networks. Furthermore, the number of nodes of a particular sub-network might be restricted.

In principle, each location must be equipped appropriately according to its hierarchy level. This includes the infrastructure (building, equipment footprint, electricity, air conditioning, etc.), detailed equipment (switches, routers, line and tributary cards, interfaces, backup, accounting, etc.) as well as maintenance and operations personnel. It is also possible to consider these requirements at a very accurate level, however, given the imprecision of the planning data it is not useful to go that much into detail. Often the planning horizon is too long to obtain an accurate traffic forecast, to obtain precise equipment cost, or even the technological specifications of future equipment. Therefore, it is advisable to use known information about the network locations in aggregated form, that is, to consider for each location a small and fixed list of admissible configurations, and to select a configuration for each location in a solution.

There are also alternatives to connect the different network locations. What type of link is available to connect two particular locations typically depends on their levels in the hierarchy and their potential configurations. For the G-WiN, for instance, only between backbone locations it was possible to install high-speed connections of 2.4 Gbit/s or 10 Gbit/s. As additional restriction, the distance in kilometers of a connection might be restricted, or in countries with high mountains and long valleys, many connections might not even be possible, although the distance in kilometer is relatively small. Similar to the locations, it is therefore useful to consider a small and fixed list of admissible configurations for each possible connection.

The access networks are supposed to be strongly hierarchical, that is, each location of a lower level is connected to *exactly* one location of the next higher level.

2.2 Routing and survivability

The access networks have a star- or tree-like structure, and therefore, the routing of communication demands follows from the definition of the hierarchy (see bolder lines in Figure 1). Only in the backbone the routing depends on the particular routing protocol. In any case, the capacities selected in the all sub-networks must be sufficient to permit a feasible routing.

The OSPF routing protocol is used in the G-WiN. Therefore, some peculiarities must be taken into account when dealing with the routing in the backbone. First, in order to guarantee a non-ambiguous, non-bifurcated shortest path routing, the routing weights are required to induce *unique* shortest paths between all pairs of backbone locations. If the routing weights do not satisfy this requirement but non-bifurcated routing is enforced, it solely depends on the implementation of the routing protocol in the operational network, which of the shortest paths is chosen in case of ambiguity. As the administrator may loose some control over the routing and capacities might no longer be adequate for the real traffic flows, such routing weights should be strictly avoided. Second, the weights in routing protocols such as OSPF are bounded. For OSPF, for instance, the weights must be integer numbers in the range between 0 and 65535. If one is not interested in minimizing the maximum or the sum of the weights, however, but only in finding integer weights within the given bounds, a simple scaling-and-rounding procedure or solving the integer inverse shortest path problem (see [3]) for the path set induced by the given weights are computationally admissible for real-world size networks.

2.3 Optimization target

The objective is to identify a network structure which minimizes the overall cost, comprising the cost of the equipment, the locations, the personnel as well as the connection costs.

3 Mathematical model

In this section, the integer programming model used to describe the planning task in terms of parameters, variables, inequalities and an objective function is introduced.

3.1 Topology, hardware and link capacities

Given are all potential node locations and point-to-point links which may be included in the final topology. This potential network is modeled by an undirected *supply graph* $G = (V, E)$, where V is the set of node locations and E is the set of links. The subset of all potential backbone nodes is denoted by $W \subseteq V$. All nodes in $V \setminus W$ are access nodes. Links between access nodes are not possible, that is, for all $uv \in E$ either $u \in W$ and/or $v \in W$. Furthermore, M is the maximal number of access nodes that can be served by a backbone node. For each undirected link $uv \in E$, we denote by (u, v) and (v, u) its two associated directed arcs and set $A := \{(u, v), (v, u) : uv \in E\}$.

The list T_W specifies the potential *node configurations* with capacities. Similarly, T_B denotes all potential *link configurations* between backbone nodes and T_A all potential link configurations between access and backbone nodes. The set of all potential link configurations is $T := T_A \cup T_B$, and for a particular link $uv \in E$, $T_{uv} := T_A \cup T_B$ if $u, v \in W$, and $T_{uv} := T_A$ otherwise. Each node configuration $t \in T_W$ provides a capacity c^t and each link configuration t of an undirected link $uv \in E$ provides a bidirected routing capacity c^t . The node capacity must be larger than the sum of capacities of all attached links.

To model the decision whether a location is backbone or access node, we introduce variables $x_v^t \in \{0, 1\}$ for all potential backbone nodes $v \in W$ and all available node configurations $t \in T_W$. Notice that v is selected as a backbone node if inequality (2) holds at equality. To model both the selection of connections between the access and the backbone network and the particular link configuration, we introduce the variables $x_{uv}^t \in \{0, 1\}$ for all $uv \in E$ and all $t \in T_{uv}$, technology installed on uv .

$$\sum_{t \in T_A} x_{uv}^t + 2 \sum_{t \in T_B} x_{uv}^t \leq \sum_{t \in T_W} (x_u^t + x_v^t) \quad \begin{array}{l} u, v \in W \\ uv \in E \end{array} \quad (1)$$

$$\sum_{t \in T_W} x_v^t \leq 1 \quad v \in W \quad (2)$$

$$\sum_{t \in T_A} x_{uv}^t \leq \sum_{t \in T_W} x_v^t \quad \begin{array}{l} v \in W \\ u \in V \setminus W \\ uv \in E \end{array} \quad (3)$$

$$\sum_{u \in \delta(v)} \sum_{t \in T_A} c^t x_{uv}^t \leq \sum_{t \in T_W} c^t x_v^t \quad v \in W \quad (4)$$

$$2 - \sum_{t \in T_A} x_{uv}^t \geq \sum_{t \in T_W} (x_u^t + x_v^t) \quad \begin{array}{l} u, v \in W \\ uv \in E \end{array} \quad (5)$$

$$\sum_{uv \in \delta(v)} \sum_{t \in T_A} x_{uv}^t \geq 1 - \sum_{t \in T_W} x_v^t \quad v \in W \quad (6)$$

$$\sum_{uv \in \delta(v)} \sum_{t \in T_A} x_{uv}^t \leq 1 + M \sum_{t \in T_W} x_v^t \quad v \in W \quad (7)$$

$$\sum_{uv \in \delta(v)} \sum_{t \in T_A} x_{uv}^t = 1 \quad v \in V \setminus W \quad (8)$$

Inequalities (1) and (3) ensure that for each installed access link at least one and for each installed backbone link both terminal nodes are indeed chosen as backbone nodes. Inequality (4) ensures the compatibility of link and node configurations. Inequality (5) ensures that access link technologies cannot be

chosen between backbone nodes. Inequalities (6), (7) and (8) ensure that each access node is connected to exactly one backbone node.

Notice that inequalities (1)–(8) only guarantee that each access node is connected to exactly one backbone node and that the chosen link technologies match the type of the terminal nodes. This system does not guarantee that the (backbone-) network is connected. Appropriate connectivity is enforced by the flows via the flow formulation in the next section and via additional metric inequalities [12].

Let \mathcal{C} be the space of all node and link configuration variables introduced above. We define $X := \{x \in \{0, 1\}^{\mathcal{C}} : x \text{ satisfies (1)–(8)}\}$.

3.2 Routing

For the solution approach presented in Section 4, explicit variables describing the routing paths of the demands are not needed. The following model of a non-bifurcated shortest path routing implicitly describes the relation between the routing weights and the resulting link flows.

For each pair of nodes $u, v \in V$, let $d^{u,v} \in \mathbb{R}_+$, denote the directed traffic demand from u to v . The routing weights assigned to the arcs $(u, v) \in A$ are modeled by the variables $w_{(u,v)} \in \mathbb{N}$. The traffic flows resulting from the chosen routing weights for the given demands are expressed by the edge flow variables $f_{(u,v)} \in \mathbb{R}_+$, $(u, v) \in A$. We define $F := \{(\mathbf{f}, \mathbf{w}) \in (\mathbb{R}^A, \mathbb{N}^E) : \mathbf{w} \text{ induces unique shortest paths, } \mathbf{f} \text{ is induced flow wrt. demands } \mathbf{d}\}$.

With this notation, the problem of finding a shortest path routing can be formulated as follows:

$$(\mathbf{f}, \mathbf{w}) \in F, \tag{9}$$

$$f_{(u,v)} \leq \sum_{t \in T_{uv}} c^t x_{uv}^t \quad \forall (u, v) \in A. \tag{10}$$

Inequality (9), which involves the implicitly defined set F , ensures that \mathbf{w} induces unique shortest paths for all demands and that \mathbf{f} is the corresponding flow on the directed links. Inequality (10) guarantees that this flow does not exceed the provided link capacities.

3.3 Cost minimization

The objective of the network design problem is to minimize the total network cost: These comprise the cost k_v^t of setting up a backbone node at $v \in W$ with configuration t and the cost k_{uv}^t of installing link configuration t between u and v :

$$\min \mathbf{k}^T x := \sum_{v \in W} \sum_{t \in T_W} k_v^t x_v^t + \sum_{uv \in E} \sum_{t \in T_{uv}} k_{uv}^t x_{uv}^t. \tag{11}$$

4 Algorithmic Approach

In this section, we present our solution approach for the integrated network design and shortest path routing problem and discuss some implementation details.

The complete model (1)–(11) can be written as

$$\left\{ \begin{array}{l} \min \mathbf{k}^T \mathbf{x} \\ \text{s.t. } f_{(u,v)} \leq \sum_{t \in T_{uv}} c^t x_{uv}^t \quad \text{for all } (u,v) \in A \\ \mathbf{x} \in X \\ (\mathbf{f}, \mathbf{w}) \in F \end{array} \right\} \quad (12)$$

Relaxing the capacity constraints, let $\mu_{(u,v)} \geq 0$, for all $(u,v) \in A$, be the associated Lagrangian dual multipliers. Then, the resulting Lagrangian function is

$$\begin{aligned} L(\boldsymbol{\mu}) &= L^X(\boldsymbol{\mu}) + L^F(\boldsymbol{\mu}) \\ &:= \min \left\{ \mathbf{k}^T \mathbf{x} - \sum_{(u,v) \in A} \sum_{t \in T_{uv}} \mu_{(u,v)} c^t x_{uv}^t : \mathbf{x} \in X \right\} \\ &\quad + \min \left\{ \sum_{(u,v) \in A} \mu_{(u,v)} f_{(u,v)} : (\mathbf{f}, \mathbf{w}) \in F \right\} \end{aligned}$$

For each dual vector $\boldsymbol{\mu} \in \mathbb{R}_+^A$, the value $L(\boldsymbol{\mu})$ is a lower bound for the optimal value of the original problem (12). Hence,

$$L^* := \max_{\boldsymbol{\mu} \in \mathbb{R}_+^A} L(\boldsymbol{\mu}) \leq \mathbf{k}^T \mathbf{x}^*,$$

where $(\mathbf{x}^*, \mathbf{f}^*, \mathbf{w}^*)$ is an optimal solution of (12). As there are only finitely many different (basic) solutions \mathbf{x} and \mathbf{f} with $\mathbf{x} \in X$ and $(\mathbf{f}, \mathbf{w}) \in F$, both functions $L^X(\boldsymbol{\mu})$ and $L^F(\boldsymbol{\mu})$ are concave in $\boldsymbol{\mu}$. Hence, $-L(\boldsymbol{\mu})$ is convex and the problem of finding the optimal dual multipliers $\boldsymbol{\mu}$ can be solved by a general convex function optimization algorithm.

Notice that exactly the same Lagrangian function $L(\boldsymbol{\mu})$ and the same lower bound L^* is obtained, if the analogous network design problem with general multicommodity flow routing instead of non-bifurcated shortest path routing is relaxed. Hence, we cannot expect to obtain very tight lower bounds with this approach.

Nevertheless, this Lagrangian approach is attractive for practical computations. One reason is that the Lagrangian function $L(\boldsymbol{\mu})$ decomposes into the sum of two functions $L^X(\boldsymbol{\mu})$ and $L^F(\boldsymbol{\mu})$, both of which can be evaluated efficiently for real-world size networks.

Evaluating the first function $L^X(\boldsymbol{\mu})$ corresponds to the problem of finding a valid network structure and hardware installation that minimizes a linear objective function. Although this problem may be \mathcal{NP} -hard in general, its integer programming formulation can be solved very efficiently by state-of-the-art integer programming solvers for real-world problems. Notice that this formulation contains only the variables \mathbf{x} and the inequalities (1)–(8). The traffic demands and flows only affect the objective function coefficients via the Lagrangian dual multipliers. For simple hardware models, this problem decomposes even further. For example, if only one capacity may be chosen for each link, choosing for each edge the minimum cost capacity yields a cost minimal overall hardware installation.

The optimization problem associated with evaluating the second function $L^F(\boldsymbol{\mu})$ can be solved by any shortest path algorithm. The task is to find a non-bifurcated shortest path routing that minimizes the total flow costs for $\boldsymbol{\mu}$. It is not hard to see that choosing for each pair of nodes the shortest path with




Figure 2: Potential network

respect to μ yields the optimal solution. Ties between equally long shortest paths can be broken by using an arbitrary numbering of the nodes or links as a secondary length function; similar to perturbing μ accordingly, see [4]. This guarantees the existence of some integer routing weights that induce the same set of shortest paths as μ .

Another reason for this Lagrangian relaxation approach being computationally interesting is the possibility to include other heuristics. After each iteration of the convex optimization algorithm, the current duals μ can be easily interpreted as routing weights, as in the evaluation of $L^F(\mu)$. In practice, these weights seem to provide good starting points for other heuristics that are based on evaluation of routing weights. From this perspective, our Lagrangian relaxation approach also can be seen as a primal heuristic that modifies the current routing weights according to the dual information and, as a byproduct, produces a lower bound for the optimum solution value.

5 Computational results

In this section, we describe results of applying the algorithms described in Section 4 to a planning scenario we have been investigating for the German Research Network (G-WiN), which is operated by DFN-Verein.

5.1 Planning scenario

The G-WiN comprises 759 locations, which needed to be organized in a 2-level hierarchy: backbone and access. For survivability reasons, the backbone is required to be 2-node connected, and each access location is attached to exactly one backbone location. There are 30 potential backbone nodes, but for operational reasons at most 20 of them can be selected as backbone location. For each access node, 10 links are introduced to define the potential connections to the backbone. These are pre-selected according to geographical criteria such as the distance in kilometer.

At all potential backbone locations, either an access router or a Cisco R1200 can be installed. In terms of optimization, the cost for access routers is unavoidable and can therefore be neglected. It is only necessary to account for the additional cost incurred by selecting a node as a backbone node: this cost is

1000000 for each R1200.

The available link capacities to connect access to backbone routers, or to connect two backbone routers are:

access–backbone: 128K, 2M, 34M, 155M, 622M
backbone–backbone: 155M, 622M, 2.4G, 10G

The cost associated with installing some of these capacities on the available links are:

capacity	fix-cost	km-length	km-cost
128K	3278	15-50-999	131-32-19
2M	6093	15-50-999	401-111-53
34M	45691	15-50-999	1523-623-429
155M	52015	15-50-999	2109-1054-436
622M	104030	15-50-999	4218-2108-872
2400M	208060	15-50-999	8436-4216-1744
10000M	1000000	15-50-999	17000-8500-4000

Table 1: Link capacities and cost.

For example, installing a link capacity of 2 Mbit/s incurs fixed cost 6093 and additional cost depending on the length of the link: Each kilometer up to 15 km costs 401, each kilometer between 15 and 50 costs 111 and each kilometer above 50 costs 53.

The DFN continuously performs traffic measurements in the network and has therefore means to provide accurate traffic distributions. As indicated in Figure 3, the structure of the communication demands is very heterogeneous and only very few locations handle large fractions of the overall traffic.

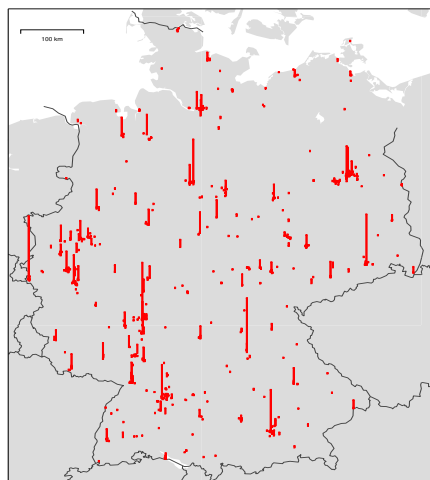


Figure 3: Source-demands




Figure 4: Best known solution

5.2 The results

To solve this planning task, the Lagrangian relaxation approach presented in Section 4 is applied. The implementation is part of the DISCNET network optimization library [2]. The data structures and algorithms are based on the standard C++ library and LEDA [1]. The CONICBUNDLE algorithm of Helberg [8, 9] is used to solve the convex optimization problems. All linear and integer programs are solved by CPLEX [11].

With this algorithmic framework, we have computed solutions within a time frame of 15 minutes on a typical PC. The best solution is shown in Figure 4. The cost of this solution is 59680665.

This is a solution to a particular planning scenario which satisfies *all* planning requirements. However, given that an operator has a tool or an algorithm which computes feasible network configurations like this, there is still a number of open questions to answer:

- What is the quality of this solution?
- Could another tool compute a solution that is 20 percent better?
- Is it worth spending more time to find a better solution?

As already mentioned, the Lagrangian approach presented in Section 4 provides, as byproduct, a lower bound for the optimum solution value. Using this lower bound, it is possible to seriously answer the questions stated above. For the planning instance described above, a lower bound of value 59407764.1 has been computed and therefore, there is an optimality gap of 0.459 percent. In other words, it might be that an optimal solution has been computed by our algorithm, but even in the worst case, any other solution methodology can be at most 0.459 percent better. That has inherently been proven at runtime. Consequently, operators can decide on a reliable basis and it becomes possible to make conceptual comparisons in contrast to comparisons which are based on heuristical results of undetermined quality.

As a final note, let us mention that solving such a complex problem with high accuracy is only possible due to special relaxation techniques which consider most of the variables and constraints *implicitly*. In the course of solving this planning instance, for instance, only 16.000 variables and 40.000 constraints have

been generated, while the complete mixed-integer program contains already about $8.9 \cdot 10^{30}$ different routing variables.

6 Conclusions

In this article, we presented a mathematical model and an algorithm to support one of the central strategic planning decision of network operators: How to structurally organize a large number of location in hierarchical levels? As one big advantage of our approach, not only solutions but also worst-case quality guarantees are provided. Using this we have been able to compute almost optimal solutions for a G-WiN planning instance of DFN with 759 locations.

References

- [1] Algorithmic Solutions Software GmbH, Schützenstr. 3-5, D- 66123 Saarbrücken, Germany. LEDA—*Library of Efficient Data types and Algorithms*, 1998–2003. Information available at <http://www.algorithmic-solutions.com>.
- [2] atesio GmbH, Rubensstr. 126, D-12157 Berlin, Germany. DISCNET, 2000–2003. Information available at <http://www.atesio.de>.
- [3] W. Ben-Ameur and E. Gourdin. Internet routing and related topology issues. submitted to SIAM J. Discrete Mathematics, 2000.
- [4] A. Bley, M. Grötschel, and R. Wessäly. Design of broadband virtual private networks: Model and heuristics for the B-WiN. In N. Dean, D. F. Hsu, and R. Ravi, editors, *Robust Communication Networks: Interconnection and Survivability*, volume 53 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 1–16. AMS, 1998.
- [5] A. Bley and T. Koch. Optimierung des G-WiN. DFN-Mitteilungen, 54:13–15, 2000.
- [6] A. Bley and T. Koch. Integer programming approaches to access and backbone IP-network planning. Technical Report ZR 02-41, Konrad-Zuse Zentrum für Informationstechnik Berlin (ZIB), 2002.
- [7] DFN-Verein zur Förderung eines Deutschen Forschungsnetzes. Information available at <http://www.dfn.de>.
- [8] C. Helmberg and K.C. Kiwiel. A spectral bundle method with bounds. *Mathematical Programming*, 93(2):173–194, 2002.
- [9] C. Helmberg and F. Rendl. A spectral bundle method for semidefinite programming. *SIAM Journal on Optimization*, 10(3):673–696, 2000.
- [10] T. Koch and R. Wessäly. Studie zur Optimierung des MSC-Backbone von E-Plus. Internal Document (confidential), Konrad-Zuse Zentrum für Informationstechnik Berlin (ZIB), 2000.
- [11] ILOG CPLEX Division, 889 Alder Avenue, Suite 200, Incline Village, NV 89451, USA. *ILOG CPLEX 8.0 Reference Manual*, 2002. Information available at <http://www.cplex.com>.
- [12] R. Wessäly. *Dimensioning Survivable Capacitated Networks*. PhD thesis, Technical University Berlin, Germany, 2000.