Robustness in Communication Networks: Scenarios and Mathematical Approaches

Andreas Bley(1), Fabio D'Andreagiovanni(2), Andreas Hanemann(3)

(1)Institut für Mathematik, Technische Universität Berlin, Straße des 17. Juni 136, 10623 Berlin, bley@math.tu-berlin.de
(2)Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), Takustraße 7, 14195 Berlin, d.andreagiovanni@zib.de
(3)DFN-Verein, Alexanderplatz 1, 10178 Berlin, hanemann@dfn.de

Abstract

The planning of wide-area networks to achieve robust operation is an ongoing challenge for network providers. Robust means to ensure stable operation of a network in case of fault occurrence. In this paper several scenarios in the context of the German Research Network are discussed to point out robustness aspects in a more detailed manner. The approach to address these challenges is to apply mathematical methods, more precisely exact and approximate integer linear programming and fast combinatorial algorithms.

1 Introduction

Wide-area networks are composed of several network layers which depend on each other in a complex manner. Therefore, the analysis of the consequences of faults is a challenge as well as the methods to be applied for making a network robust. Robustness refers to having a stable operation of the network even in fault situations (e.g. when fiber cuts happen). The planning to construct a robust network gets very complex once the topology of the network allows more and more different routes. Therefore, mathematical methods to address the challenges are discussed in this paper, which reports on parts of the ROBUKOM (Robust Communication Networks) project [ROBUKOM]. This project has been started in October 2010 and is funded by the German Federal Ministry of Education and Research (BMBF). DFN (German Research Network) contributes to the project by providing real world scenarios and by evaluating the prototypical implementations of network optimization tools.

The paper is organized as follows. In Section 2 DFN's network (called X-WiN) is presented as a reference example of a state-of-the-art data network and scenarios for making the network more robust are being discussed. The mathematical approaches adopted in the ROBUKOM project to address the scenarios are explained in Section 3. The section provides an overview of related work as well. Finally, Section 4 concludes the paper.

2 Network Optimization Scenarios

The network robustness scenarios which are outlined in the following are set in context with DFN's network, which serves as national research and education network to connect German universities and research institutions.

DFN's network is based on a dark fiber infrastructure (see Figure 1). There are only a few places remaining in the network where single wavelengths have been leased. On the fiber infrastructure wavelengths are being provided based on the DWDM (Dense Wavelength Division Multiplexing) technology. DWDM technology has been installed at all of the 58 points-of-presence (PoPs) of the network.

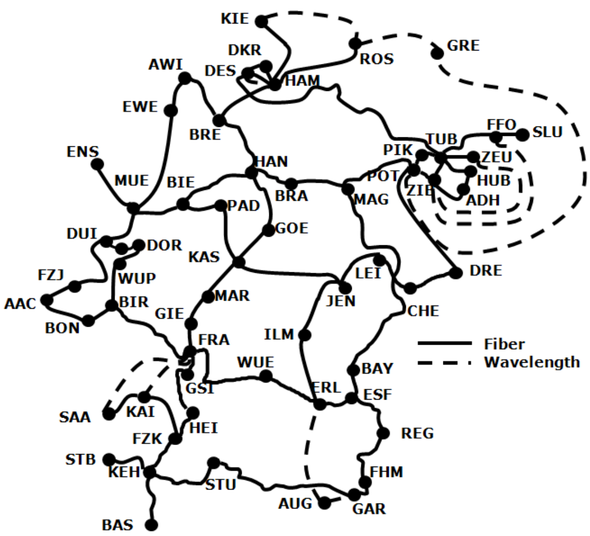

Figure 1 Fiber topology of DFN's network (February 2011)

The network is a so-called hybrid network, because the wavelengths are used for two purposes. One of the purposes is to connect routers which are installed at 51 network locations (see Figure 2). The interconnections of the routers form an IP network providing the so-called DFNInternet Service. This is the most important service which is relevant for all users of DFN.

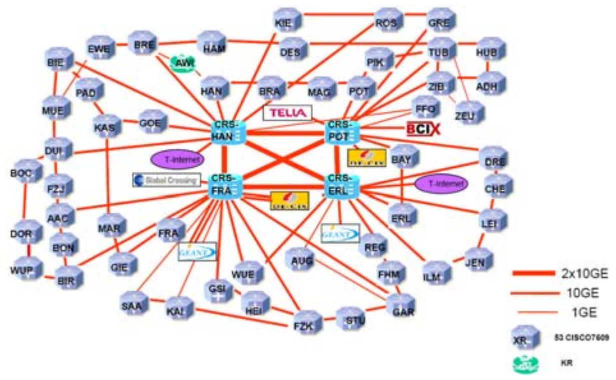

Figure 2 Router topology of DFN's network (February 2011)

This IP Service is depicted in Figure 3 in a simplified manner. Four PoPs are shown where routers and DWDM equipment are located. These are interconnected on different levels: with physical fibers shown as solid lines, with wavelengths being transmitted inside the fibers shown as dashed lines, and with logical IP links being realized via dedicated wavelengths shown as dotted lines. In the figure, the logical IP links are mostly realized just to the neighboring PoPs. But there is also one IP link from the PoP on the left going to the PoP on the right which is realized using a wavelength that passes through the PoP at the bottom.

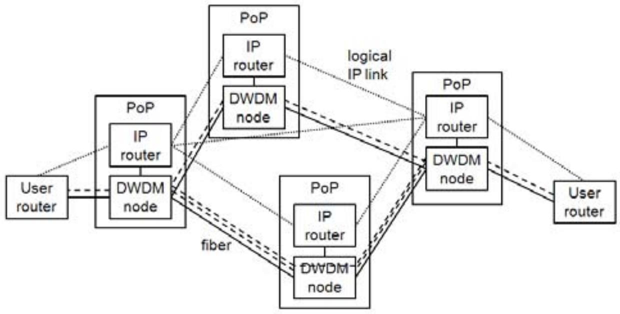

Figure 3 Simplified technical view on the IP Service

The second purpose of the wavelengths in DFN's network is to connect user sites where high volumes of data are being transferred. In this context it is reasonable not to apply IP routing inside of DFN's network, but to dedicate certain wavelengths specifically to these connections. They are called Layer 2 VPNs (Virtual Private Networks).

A view onto the L2 VPN Service is given in Figure 4. The IP routers in the network are not shown because only the DWDM equipment is needed to realize the service. In the example shown in the figure two wavelengths are used to realize an optical 1+1 protection inside the network for the end-to-end user connection (one time passing the upper PoP and the other time passing the lower PoP). A user configured end-to-end IP connection is not shown in the figure.

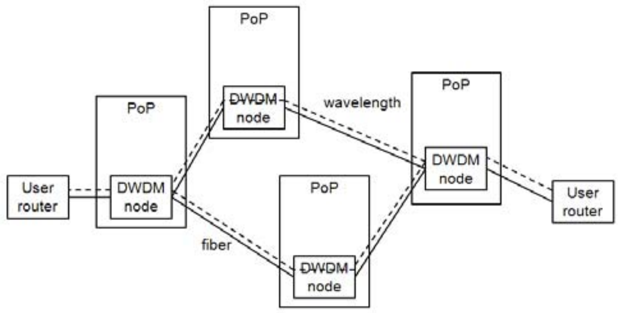

Figure 4 Simplified view on the L2VPN Service

To ensure network robustness DFN already applies several measures. The DWDM equipment as well as the routers have a high degree of internal redundancy and the power supply at the PoPs is ensured with means such as uninterrupted power supply units. Every user (e.g. a university) is supported by DFN to realize a redundant connection to DFN's network. For the IP services, DFN does not use optical 1+1 protection of the links, because the rerouting mechanisms of IP are applied. Here at least two, but often three or more paths, are available and it is ensured that these paths are actually realized via diverse fiber paths. For the Layer 2 VPNs it is not possible for DFN to use IP protection mechanisms. Therefore, optical 1+1 protection mechanisms are applied so that a preconfigured backup path is available for automatic failover protection.

Based on all these measures DFN is already achieving very high availability. Nevertheless, investigations about additional measures are highly relevant because failures of Internet access can today endanger the proper functioning of complete institutions.

In the following the challenge of making a network more robust is addressed on the basis of analyzing real world scenarios.

2.1 Additional wavelengths scenario

In the first scenario, the IP service is considered. The question is whether in a given network it is reasonable to install some additional wavelengths for IP links between routers in order to realize a robustness gain. In this scenario all wavelengths have a capacity of 10 Gbit/s.

To address this question the meaning of "robustness gain" has to be investigated. It is strongly related to the definition of robustness. Here the idea is to define this dependent on the utilization in the network which should be as low as possible for the maximum found among all links given a certain traffic demand. The motivation for using the utilization as robustness measure is that high utilization leads to high delay and jitter or even to packet loss which should be avoided. The utilization is not only considered in the basic situation without the occurrence of faults, but also in all situations with a single fiber cut and in all situations with a single node failure. The overall robustness is then derived by combining all scenarios.

Based on a given matrix of existing traffic demand, which shows how much traffic is exchanged between each pair of PoPs in the network, it should be investigated where it is beneficial to install additional IP links. As such an additional IP link is realized here via a dedicated wavelength the scenario is called “additional wavelengths scenario”. The challenging aspect of the task is therefore to cope with the amount of possibilities for placing such a link into the network.

Finally, the investigation of the scenario should lead to a decision aid where the robustness gain is compared to the number of additional wavelengths. A potential conclusion can then be to say that a certain number of additional wavelengths is useful, while more wavelengths lead only to an irrelevant robustness gain.

A remark on the scenario is that it is based on the currently deployed technology, where transponders are specifically manufactured for a given path in the network. Therefore, only additional wavelengths are considered and the possibility to change existing network paths is not taken into account. However, recent DWDM systems have become much more flexible to measure the fiber characteristics along a given path and to tune the lasers accordingly. Relocation of equipment is thus possible and, as a consequence, changes to existing wavelengths can be considered as well.

2.2 Additional fibers scenario

A related scenario is based on the possibility to rent additional fibers. Here the idea is not to consider all possible additions of fibers to the network, but to assume a concrete offer for one or few additional fibers. The question in such a situation is whether the inclusion of these additional fibers is reasonable from a robustness point-of-view.

In this scenario the more advanced DWDM hardware is assumed so that also existing wavelengths can make use of an additional fiber. Otherwise, the additional fiber would require additional wavelengths in any case.

2.3 Introduction of 100 Gbit/s scenario

Wide-area networks of today are based on 10 Gbit/s links, but 40 Gbit/s and 100 Gbit/s wavelengths and router interface cards are already available or are about to be made available in the near future. The question that arises is how to introduce such wavelengths into a network in a reasonable manner from a robustness point of view. Single 100 Gbit/s connections are likely to become critical parts of the network, in particular if they replace multiple 10 Gbit/s connections. It therefore has to be investigated which backup installations are necessary to cope with potential failures of 100 Gbit/s wavelengths.

2.4 Migration scenario

Since DWDM technology has significantly improved over the last years, many networks would greatly benefit from migrating to a new network generation. Such a migration is complicated from a robustness point of view because there will be intermediate situations where a part of the network has already been migrated to a new DWDM generation, while some other parts are still based on the old equipment. This means that protection mechanisms will not work for certain periods of time, thus entailing issues especially for L2 VPNs. It is therefore desirable to have a decision support tool which helps to elaborate an order of migration steps to be carried out.

3 Mathematical Modeling of Scenarios

As highlighted in Section 1, wide-area networks consist of a complex stack of subnetworks that are based on different technologies and are commonly called *layers*. The stack is hierarchically ordered and adjacent layers operate according to a client-server relation.

Designing a multi-layer network essentially consists in dimensioning the two main elements of every layer, nodes and links, in order to route traffic demand generated by users, while satisfying a set of side constraints. In particular, the nodes must be equipped with proper hardware, capacities must be installed on the links, and the routing of the traffic demands must be decided.

Traditionally, the design task has been operated by a sequential top-down approach: starting from the top of the stack, each layer is individually designed and the output configuration is used as input for the design of the lower level. This sequential approach appears attractive to network planners because of its lower complexity and higher computational tractability. However, it neglects the strong interdependencies between layers and thus may lead to non-optimal solutions. In planning scenarios where the network robustness is of crucial importance, it is essential to consider the interdependencies between the different network layers and the integrated design of the layers becomes essential.

In the ROBUKOM project we therefore consider the integrated design of a two-layer network made up of a *physical layer* corresponding to the optical fiber network and a *logical layer* corresponding to the wavelengths between the IP routers or the nodes of the L2 VPNs.

In the following sections, we sketch the mathematical models and solution approaches used to address the first three planning scenarios introduced in Section 2. These are all related to the robustness of the IP Service network. The migration scenario, which is related to the replacement of the DWDM technology, has a different nature and is associated with a different statement of the problem and thus a different optimization model. Therefore, this scenario is separately discussed in Section 3.3.

3.1 The IP network scenarios

In general, the problems that we consider in the first three scenarios can be described in the following way:

The physical and the logical network are modeled as two undirected graphs, denoted by G(V,E) and H(V,L), respectively. The physical network models the fiber network. It also includes extra links for the physically disjoint routes of the few remaining leased wavelengths. Each physical link $e \in E$ has an associated length expressing its fiber length (in km). The logical network models the wavelength connections. Each logical link $l \in L$ corresponds to an existing or potential new wavelength established in the physical graph H(V,L) or to an existing leased wavelength. As a consequence, a pair of logical nodes $i,j \in V$ may be connected by multiple parallel logical links that correspond to different paths in the physical network. Note that not all links of the logical network are given explicitly. As new wavelength connections can be set up between (almost) all pairs of nodes, only the logical links corresponding to the wavelengths of the existing network are given explicitly in general. The logical links that correspond to potential new wavelengths are given implicitly as all those logical links, where the length of the physical transmission path does not exceed the maximum reach of the used transceiver types, which is given with the problem input.

In the additional wavelengths scenario and in the 100 Gbit/s scenario, the links of the physical network describe the topology of the existing fiber network (together with additional links for the leased wavelengths), while the explicitly given logical links describe the existing wavelengths. In the additional fibers scenario, the physical network also includes links (and maybe further nodes) for those fibers, whose addition to the existing fiber network shall be evaluated.

Every logical link has an associated bandwidth capacity. For the links corresponding to the existing IP network, this is the existing wavelengths' capacity (in Gbit/s). For the (implicitly given) optional new wavelengths, exactly one capacity from a given set of valid capacities must be chosen. In the additional wavelengths scenario and in the additional fibers scenario, only the capacity of 10 Gbit/s is considered for each new wavelength. In the 100 Gbit/s planning scenario, we additionally specify a subset U of the nodes as 100 Gbit/s enabled. For those new logical links that have both end-nodes in U, a capacity of either 10 Gbit/s or 100 Gbit/s may be chosen. For all other logical links, only the 10 Gbit/s capacity is considered.

In a more general IP-over-DWDM network design problem setting, we also may consider the setup costs and channel capacities of the fibers and the modular IP and optical hardware installed at the nodes of the network [COST293, Or09]. However, as the detailed hardware installation and configuration is not of primary interest in the described planning scenarios, we do not model these aspects here.

The IP traffic generated by users is modeled as a set K of directed point-to-point communication demands. Each demand $k \in K$ is associated with a source node $s^k \in V$, a target node $t^k \in V$, and a traffic demand $d_k \in \mathbb{R}^+$ that must be routed from s^k and t^k. In our planning scenarios, the demands are given as a single (deterministic) traffic matrix, whose entries correspond to the end-to-end bitrates that shall be realizable by routing all demands simultaneously. The demand values are traffic forecasts derived from traffic measurements in the current network.

In the IP Service network of DFN, all traffic demands are routed using the OSPF (Open Shortest Path First) routing protocol unsplit (i.e., via single paths) from their sources to their destinations. In case of a node or a link failure, the OSPF restoration mechanism is used to re-compute the shortest paths in the residual network and re-route the traffic demands accordingly.

To model this behavior appropriately in our optimization model, we consider a set of network operating states F. This set consists of the normal state, where all links and nodes are operational, all single physical link failures, that model the cases where a single fiber and all wavelengths using that fiber fail, and all single node failures, that correspond to the cases where an entire node (including all its adjacent fibers) fail. Note that in the latter cases, the single node failures, also all the demands starting or ending at the failing node disappear. The routing of the traffic demands within the logical network is modeled as a unique shortest path routing. The routing weights, which are used to compute the routing paths in the normal state and in the failure states, are chosen by the optimization. Note that the routing weights must be chosen in such a way, that they define a unique shortest (s_r,t_r)-path for every (non-disappearing) demand in every operating state. Load balancing among parallel logical links is allowed, but multi-path routing via different paths in the logical network (as in Equal Cost Multipath Routing) is not permitted.

The chosen logical links and routing weights define a vector of directed traffic flows on the logical links for each operating state. Together with the existing and newly chosen logical link capacities, this defines a vector of link utilization values for each operating state. In order to assess the robustness of the network, we consider three key figures: the maximum utilization over all logical links in the normal network state, the maximum utilization over all logical links in any single physical link failure state, and the maximum utilization over all logical links in any single node failure. Given both hard upper bounds and soft averaging factors for each of these three key numbers, the goal is to minimize the weighted average of the key utilization numbers, while respecting their hard upper limits.

In each of the three considered planning scenarios, two kinds of decisions must be taken: 1) the choice and the capacities of the logical links to be added to the existing network and 2) the routing of the traffic demands within the resulting logical network. Note that the capacity of the new logical links is fixed to 10 Gbit/s in the additional wavelengths scenario and in the additional fibers scenario. In these two planning scenarios, the output of the optimization consists only of the set of chosen new logical links (together with their corresponding paths in the underlying physical network) and the OSPF routing weights for the logical links. Only in the 100 Gbit/s scenario, where different logical link capacities are considered, we have to actually decide upon the capacity assigned to the new logical links. So, in this planning scenario, the optimization output includes the new logical link capacities in addition to the set of fiber paths of the new logical links and the routing weights.

In order to cope with the fact that the existing network cannot be changed completely within a single reconfiguration step in practice, we consider versions of the above optimization problems, where the number of logical links added to the existing network is bounded by a parameter given as problem input. By solving the optimization problems for different numbers of permitted new logical links, we obtain so-called Pareto curves describing the relation between the number of logical links to be added to the network and the gain in network robustness (i.e., utilization) that can be achieved. These curves are a valuable decision support in practice: for small numbers of additional logical links, we see how the network robustness can be improved with only few additions within a single reconfiguration step, while for large numbers of additional logical links, the robustness gain indicates the potential network robustness of the underlying fiber network that could be achieved in the long run.

3.2 Solution approaches

The optimization problems described above can be modeled as Mixed-Integer Linear Programming problems (MILP) [BeTs97, NeWo88], where (binary) integer variables are used to describe the choice and the capacity of the logical links, integer variables are used to model the OSPF routing weights assigned to the logical links, and a combination of binary and/or non-negative continuous variables is used to model traffic flows induced by the routing weights in the considered operating states and the resulting utilization values. Linear equalities and inequalities on these variables describe the technical and operational side constraints that must be satisfied.

In the following, we illustrate how to model a simplified version of the additional wavelengths scenario problem, where each traffic demand must be routed unsplit along a single path in each operating state, but these routing paths need not necessarily fulfill the shortest path property that is implied by the use of the OSPF routing protocol. A computational efficient way to incorporate this property into the model is described in [Bl11]. For the sake of notational simplicity, in this model we only consider single link failures, but no node failures.

Network topology. The physical and the logical network are modeled by two undirected graphs, denoted by G(V,E) and H(V,L), respectively. The set of parallel logical links between a pair of logical nodes $i,j \in V$ is denoted by L_{ij}. Additionally, by L_{ij}^{new} we denote the set of logical links that may be added between i and j by using additional wavelengths. All the logical links of the network are assumed to have the same capacity, denoted by C, and we are allowed to activate at most N links in the set $L^{new} = \cup_{i,j \in V} L_{ij}^{new}$. By $L_e \subseteq L$ we denote the subset of logical links whose corresponding physical path contains the physical link $e \in E$.

Demands. For each directed point-to-point communication demand $k \in K$ a bandwidth d_k must be reserved in the network between s^k and t^k. We assume that each demand must be protected against the failure of single physical links $e \in E$ (in this basic illustrative model, for sake of simplicity we neglect single node failures).

Failure scenarios. We denote by F the set of all the possible failures of single physical links $e \in E$. The set includes also the scenario where no failure occurs. The link failing in scenario f is denoted by e(f). For each scenario $f \in F$, we want to define a routing plan that does not make use of the set of logical links $L_{e(f)}$ whose physical path includes the failing link e(f).

Variables. In the problem that we consider, two decisions must be taken: 1) establishing the single path on which the traffic of each demand $k \in K$ is routed in each failure scenario f∈F; 2) deciding which are the N logical links that are activated among those in L^{new}.

To model these two decisions, we introduce:

1) A set of binary variables $x_{l,ij}^{k,f}$ that indicate if the logical link $l \in L_{ij}$ is used to route the traffic of demand k∈K in the failure scenario f∈F. In particular, for each link we need two variables to distinguish the traffic flow sent from i to j from the one sent from j to i, i.e. $x_{l,ij}^{k,f}$ and $x_{l,ji}^{k,f}$;
2) A set of binary variables y_l that indicate if the logical link $l \in L^{new}$ is added to the network.

Moreover, we need two continuous variables $u_{norm} \in [0,U_{norm}]$ and $u_{fail} \in [0,U_{fail}]$, with $U_{norm}, U_{fail} \leq 1$ chosen by the network operator, to model the maximum utilization of links in the normal state and in all the failure states, respectively. If we denote by f_0 the normal state, the variables are thus defined as follows:

$$u_{norm} = \max_{i,j \in V,\, l \in L_{ij}} \frac{1}{C} \sum_{k \in K} d_k \cdot x_{l,ij}^{k,f_0}$$

$$u_{fail} = \max_{f \in F \setminus \{f_0\},\, i,j \in V,\, l \in L_{ij}} \frac{1}{C} \sum_{k \in K} d_k \cdot x_{l,ij}^{k,f}$$

Constraints. First of all, we introduce a set of constraints that model the traffic flow conservation of each demand k, in each node i of the network and in each failure scenario f:

$$\sum_{j \in V} \sum_{l \in L_{ij}} (x_{l,ij}^{k,f} - x_{l,ji}^{k,f}) = \begin{cases} 1 & \text{if } j = s^k \\ -1 & \text{if } j = t^k \\ 0 & \text{otherwise} \end{cases} \qquad i \in V, k \in K, f \in F$$

Then we need a set of constraints to model the fact that in each scenario f∈F, for every logical link $l \in L_{ij}$ between two nodes i,j∈V, the sum of all traffic flows d_k sent on l must not exceed the link capacity:

$$\sum_{k \in K} d_k \cdot x_{l,ij}^{k,f} \leq C \qquad i,j \in V, l \in L_{ij}, f \in F$$

We remark that it is necessary to introduce one capacity constraint for each of the two directions of a link $l \in L_{ij}$.
Analogously, capacity constraints must be defined also for the new logical links L^{new}. In this case, the capacity must be multiplied by the link activation variable y_l, so that no flow can be sent on l when l is not activated:

$$\sum_{k \in K} d_k \cdot x_{l,ij}^{k,f} \leq C \cdot y_l \qquad i,j \in V, l \in L_{ij}^{new}, f \in F$$

Also in this case, two capacity constraints must be defined to model the traffic flow sent on the two directions of the link.
An additional constraint is needed to express that at most N new links in L^{new} must be activated:

$$\sum_{l \in L^{new}} y_l \leq N$$

To express the fact that in failure scenario f, we may not use the logical links whose paths contain the failing link e(f), variables $x_{l,ij}^{k,f}$ such that $l \in L_{e(f)}$ must be set to zero. This actually corresponds to remove the variables from the problem.

Objective function. The aim is to minimize a weighted combination of the maximum link utilization in the normal state u_{norm} and in all failure scenarios u_{fail}, i.e.:

$$\min \quad w_{norm} \cdot u_{norm} + w_{fail} \cdot u_{fail}$$

where w_{norm}, w_{fail} are suitable weights chosen by the network operator.

Even with simpler routing models and without the consideration of multiple network failure situations, the integrated planning of two network layers is a very challenging task. Its complexity dramatically increases as soon as survivability requirements and single-path (or, even more complex, OSPF) routing requirements are included. The direct use of effective commercial Mixed-Integer Program (MIP) solvers like IBM ILOG Cplex is not practicable and alternative solution approaches must be used. Recent works have thus concentrated efforts on finding other more computationally-tractable ways of formulating the basic MIP (e.g., [KnLa07]) or on developing refined solution approaches based on exact methods (e.g., [COST293]), heuristics (e.g., [DuQiMe10]) or a mix of both (e.g., [Or09]). For a comprehensive review of related work, we refer the reader to [Or09]. Note, for example, that the integer programming model presented above contains a huge number of variables and constraints if there are many potential logical links that could - in principle – be added to the existing network. In order to solve such models computationally, one has to apply so-called column generation and cutting plane algorithms, which start with a small subset of the variables and constraints and, during the execution of the optimization algorithm, dynamically generate further variables and constraints. The advantage of these methods is to drastically reduce the size of the integer-linear programs that must be solved. This is accomplished by considering only the (small) subset of variables that is actually necessary to solve the problem [NeWo88].

Appropriately considering the shortest path restrictions implied by the OSPF routing protocol increases the computational complexity of the optimization problems even more. Finding routing weights that induce globally efficient routing paths in all network operating states is a major difficulty. Various mathematical optimization methods, including heuristics and genetic algorithms [BuReRiTh05, FoTh04, ÜmFo07], Lagrangian relaxation techniques [Bl03], and Integer-linear Programming approaches [ToPiDzZa05, Bl11] have been developed in the last decade to solve such shortest path routing optimization problems. The proposed methods, however, have some severe limitations: they are either only heuristic, i.e., they cannot guarantee to find a routing that is even close to the optimal one, or they do not scale well with the network size and the number of considered operating states. Even for pure routing optimization problems in single-layer networks, finding proven near-optimal solutions is computationally challenging if the potential node and link failures are considered. Combinations of various heuristics, mathematical decomposition techniques, and dynamic model refinement techniques based on different sub-models are necessary to control both the size and the precision of the mathematical model during the solution process. A direct integration of the OSPF routing optimization techniques with the multilayer optimization is computationally out of reach.

In all three planning scenarios, we therefore decompose the overall mathematical optimization problems into two parts. In the first part, we address the problem using a simplified routing model, where all demands must be routed in all operating states, but the routing does not have to be a valid shortest path routing. By solving this problem, we obtain a small pre-selection of those logical links (and link capacities), whose addition to the existing network leads to the biggest reduction in network utilization (at least for the simplified routing scheme), and we obtain a lower bound on the utilization reduction that can be achieved using these links. In the second part, we then solve the optimization problem using the more complex but also more precise OSPF routing model, but only allowing the logical links pre-selected in the first step to be added to the existing network. By solving this part, we decide which logical links to actually add to the existing network, which capacities to assign to these links, and how to set the OSPF routing weights.

3.3 Modeling the migration scenario

As mentioned in Section 3.1, this scenario is deeply different from the other ones and requires its own treatment. In contrast to other scenarios, it indeed presents a dynamic character: the aim is to replace the DWDM equipment everywhere in the network which is carried out in a series of steps. To avoid making the situation even more error-prone it is aimed to maintain the network topology and not to try optimizing it as part of the migration.

When setting the migration in relation to the services being offered, it can be witnessed that it poses different challenges to the IP Service and L2 VPN Service. In the case of the IP Service small disruptions have to be taken into account during the switching from old to new DWDM equipment along a fiber, but then the service would mostly have the same protection level as before. In the case of the L2 VPN Service where there are typically long backup paths which are not available during certain periods of the migration, the service is at risk. This entails that the L2 VPN Service may become unprotected and failures like fiber cuts may cause outages for a significant number of hours.

A general statement of the problem associated to the migration scenario is the following: given i) the migration horizon, ii) the network topology and services configuration, iii) the number of fibers that can be migrated at the same time (depending on how many technicians are available to be at DWDM nodes and amplifier locations at the same time), we want to define the sequence of DWDM equipment changes for the fibers. Its aim is to keep the service disruption time low, but also to minimize the risk being taken.

To our best knowledge, such migration problem has not yet been investigated and there is no previous work to which we can directly refer in literature. However, it seems reasonable to trace back the problem to a scheduling problem [NeScZi03], by introducing either integer or binary variables for each L2 VPN Service instance and DWDM equipment, to represent i) the starting period of the disconnection and reconnection operations of a service instance and ii) the starting period of the turning off of an old DWDM equipment in a node and the turning on of the new DWDM equipment in the same node.

The constraints must model the fact that a limited number of operations can be done at the same time, on the basis of the availability of technicians. Moreover, a set of constraints must be introduced to model the minimal robustness requirements: this could be done by introducing hard constraints on the maximum link utilization while the migration takes place. Additional constraints should also be introduced to model possible precedence requirements between the migration steps. Finally, the objective function (to be minimized) should correspond to a (weighted) summation of the time that service disruption happens and the time that L2 VPN connections remain unprotected.

4 Summary

In this paper several scenarios for increasing the robustness of a network have been discussed. DFN's network has served as a reference example for the discussion, but it is important to note that the topic is relevant in general. In countries where a dark fiber infrastructure is affordable, the national research and education networks are built as hybrid networks as well. In commercial networks the data transport has become very important already and a tendency towards all IP networks can be witnessed, so that at least the scenarios related to the IP Service apply for these networks as well.

To cope with the computational complexity of the mathematical optimization models arising in the presented planning scenarios, we propose a decomposition approach. It first solves a two-layer survivable network design problem based on a simplified IP routing model to compute a small subset of good candidate wavelengths to be added to the current network. In the second step, we then restrict the set of potential new wavelengths to those pre-selected in the first step and solve the combined network design and routing optimization problem with an accurate model of the IP routing protocol used in practice in order to decide which of the candidate wavelength to add to the existing network and how to set the OSPF routing weights. Note, however, that each of these two subproblems is computationally very challenging already on its own. Within the ROBUKOM project, we therefore also aim at improving the solution techniques for the two subproblems by advancing and combining heuristic and exact mathematical optimization techniques for the first subproblem and by developing dynamic and scalable model-refinement techniques for the second subproblem. This will eventually lead to more efficient solution methods for the overall planning problems described in this paper.

Acknowledgement

This paper has been supported by the German Federal Ministry of Education and Research (BMBF) as part of the ROBUKOM project (grant 03MS616E).

5 References

[BeTs97] Bertsimas, D.; Tsitsiklis, J.H.: Introduction to Linear Optimization, Athena Scientifics, New Hampshire: Nashua, 1997

[Bl03] Bley, A.: A Lagrangian approach for integrated network design and routing in IP networks. In: Proceedings of INOC 2003, pp. 107–113

[Bl11] Bley, A.: An Integer Programming Algorithm for Routing Optimization in IP Networks, Algorithmica, Vol. 60, 2011, pp. 23-45

[BuReRiTh05] Buriol, L.; Resende, M.; Ribeiro, C.; Thorup, M.: A hybrid genetic algorithm for the weight setting problem in OSPF/IS-IS routing. Networks 46, 2005, pp. 36–56

[COST293] Koster, A.M.C.A; Munoz, X.: Graph and Algorithms in Communication Networks: Springer-Verlag, Heidelberg, 2010

[DuQiMe10] Duelli, M.; Qin, X.; Menth, M.: Greedy Design of Resilient Multi-Layer Networks. Proceedings of NGI 2010. 2010.

[FoTh04] Fortz, B.; Thorup, M.: Increasing Internet capacity using local search. Comput. Optim. Appl. 29, 2004, pp. 13–48

[KnLa07] Knippel. A.; Lardeux, B.: The multi-layered network design problem. Europ. J. Oper. Res. Vol. 183, No. 1, 2007, pp. 87-99.

[NeScZi03] Neumann, K.; Schwindt, C.; Zimmermann, J.: Project Scheduling with Time Windows and Scarce Resources: Springer-Verlag, Heidelberg, 2003

[NeWo88] Wolsey, L.A.; Nemhauser, G.L.: Integer and Combinatorial Optimization. Wiley, 1988

[Or09] Orlowski, S.: Optimal Design of Survivable Multi-layer Telecommunication Networks. Ph.D. Thesis, Technische Universität Berlin, 2009

[ROBUKOM] ROBUKOM project. www.robukom.de (subprojects 3 and 4)

[ToPiDzZa05] Tomaszewski, A.; Pióro, M.; Dzida, M.; Zagozdzon, M.: Optimization of administrative weights in IP networks using the branch-and-cut approach. In: Proceedings of INOC 2005, pp. 393–400

[ÜmFo07] Ümit, H.; Fortz, B.: Fast heuristic techniques for intra-domain routing metric optimization. In: Proceedings of INOC 2007, 2007