Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ARIE KOSTER ADRIAN ZYMOLKA

Routing in optical transport networks

Optimization of the static configuration of optical transport networks

Report Milestone 4

ZIB-Report (August 2002)

ROUTING IN OPTICAL TRANSPORT NETWORKS

OPTIMIZATION OF THE STATIC CONFIGURATION OF OPTICAL TRANSPORT NETWORKS

Report Milestone 4

Arie Koster Adrian Zymolka

The work package *Optimization of the static configuration of optical transport networks* has started at September 1, 2000 and is divided into four parts each a half year long. In this paper, we report about the work and the results of the fourth and hence last project part that covers the period March 1, 2002 – August 31, 2002.

1 Introduction

The first year of the project "Optimization of the static configuration of optical transport networks" was mainly dedicated to the development of a general solution approach for the minimum cost static configuration problem of optical networks. As result, an appropriate method has been derived and implemented. This implementation serves as starting point for the second project year in which the main emphasis lies on the enhancement and in particular the application of our solution methodology.

The third half year of the project was closed by the milestone report [6]. In this report, we started discussions on the following topics:

- Integration of several survivability models.
- Continued development of advanced heuristics and exact solution approaches to achieve a minimum cost static network configuration.
- Comparison of the computed configurations according to all assessment criterions.

These topics also serve as milestones for the fourth half year, so the research on this field of tasks has been continued and extended during the ongoing project.

The structure of this report is organized as follows. Based on the prior project work (see [4], [5], and [6]), we describe in Section 2 the continued development of methods and procedures for the solution of the static optical network design problem. By their application, we have extended the computational studies. In Section 3, we discuss the performed experiments and the gained results. Finally, Section 4 concludes by summarizing the project over the complete time and gives an overview about interesting directions for further research.

2 Method advancements

In this section, we report on the advancements which have been carried out for computing minimum cost static optical network configurations. As described in [4], our solution approach bases on the decomposition of the problem into the dimensioning and routing subproblem and the wavelength assignment subproblem which are processed subsequently. Both parts are considered separately in the following description of the continued research on the applied methods and techniques.

2.1 Dimensioning and routing

The dimensioning and routing subproblem addresses two highly correlated aspects of the network configuration: the placement of optical equipment for providing sufficient capacities, and the routing of all lightpaths on top of the capacitated physical topology. In [5, Section 2], we give a detailed description of the procedure which has been developed to solve this complex task. The core of the applied method consists in the adaption of the sophisticated optimization tool DISCNET [9]. In a preprocessing step, we generate the installable capacity levels for each physical link and node as input. The application of DISCNET then yields a complete routing of all demands together with a suitably capacitated network. In addition, a lower bound on the total network cost is provided. Since the routing need not to be integer, we apply a postprocessing step which reroutes the fractional part of the demands.

Let us recall the rerouting procedure as described in [5]. First, the integer paths of the routing are fixed, while all fractional parts are removed. Note that for each demand, these fractional parts sum up to an integer number of remaining lightpaths which have to be rerouted. The lightpaths of the unsatisfied demands are then processed sequentially. The new routing of such a lightpath is determined by a shortest path computation. For this, we assign lengths to the links as follows. If a link provides so far unused channels, we set its length to zero. Otherwise, the incremental cost for installing the next capacity level is assigned as length, since the use of the link requires to extend its number of channels. After processing all lightpaths, we get an integer routing. Computational experiments have turned out that this procedure works satisfactory since usually only few fractional routings have to be rerouted.

During the third project phase, we have extended the dimensioning and routing subproblem by the development and integration of a survivability concept for the lightpath routing (see [6, Section 3.2]). This concept bases on a routing scheme which is called *diversification* and was introduced by Dahl and Stoer [2]. The use of diversification imposes additional constraints for the lightpath routing. Within each demand, the number of lightpaths which

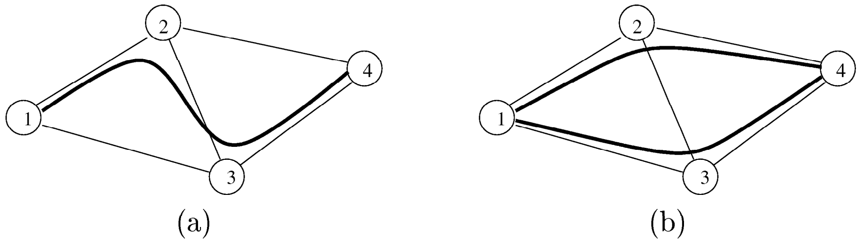

(a) (b)

Figure 1: Example for rerouting of lightpaths.

may cross the same link or node is limited. Hence, the routings of lightpaths of the same demand are not independent anymore.

As a consequence, the previously described rerouting procedure may fail as illustrated in the following small example. Consider the network in Figure 1. We assume that two lightpaths have to be routed from node 1 to node 4 with a diversification parameter of 50%, i.e., at most one of the lightpaths is allowed to cross each of the nodes 2 and 3. If all links except those between node 1 and 3 and between node 2 and 4 provide free capacity, the shortest path rerouting first chooses the lightpath depicted in the left part (a). But this blocks the second lightpath which can not be routed anymore with respect to the diversification restrictions. A feasible routing of both lightpaths is displayed in the right part (b). Note that each of these lightpaths needs to extend the capacity of a link.

To deal with the additional diversification restrictions, we have adapted the rerouting procedure. Instead of processing single lightpaths, we reroute all lightpaths which belong to the same demand in one step. For this, a minimum cost flow problem is solved in a specific network with costs and capacities. This network is constructed as follows. We split each original node into two nodes, a source copy and a target copy. A directed arc from the source copy to the target copy is inserted for each port which can be used in that node up to the limit dictated by the diversification requirements. Note that this limit may be reduced by the fixed part of the routing. In addition, we assign a cost to each of these arcs depending on the available spare ports and the costs of extending the port capacity in the original node. As many arcs as free ports are available get the cost 0. As many arcs as are available by the first extension of the port capacity in the original node get the incremental cost for this extension. This is then continued for all next extensions until all arcs have a cost assigned. Thereby, the term 'incremental cost' denotes the difference between the total cost for the extended capacity level minus the total cost for the currently installed capacity level, and hence the assigned cost value increases with each next extension considered. Next, we replace each original link by two similar sets of arcs which only differ in the connected node copies. The arcs in the first set begin at the target copy of the first node and end at the source copy of the second node, while the arcs in the second set begin at the target copy of the second node and end at the source copy of the first node. In each set, there is one arc for each channel which can be used on the link up to the limit by diversification. The costs for the arcs in each set are assigned in the same way as for the ports in nodes. As many arcs as there are spare channels on the link get the cost 0, while each capacity extension induces its incremental cost to be assigned to as many arcs as new channels are

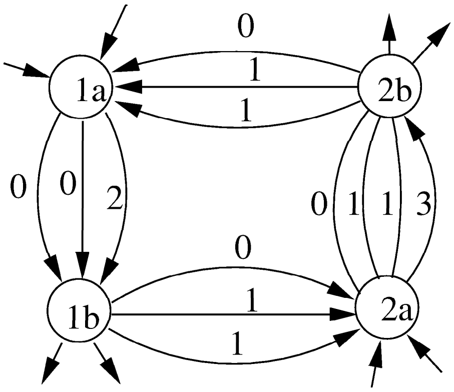

Figure 2: Example for the transformation of a linked pair of nodes for the minimum cost flow network for rerouting of lightpaths.

provided. Finally, each arc in the network has a capacity of 1.

The construction of the network is illustrated in Figure 2. The source and target copies of the original nodes 1 and 2 are indicated by the suffix a and b, respectively. At most three ports can be used in node 1, of which two are for free and the third is only available by extending the port capacity at cost 2. In node 2, four ports are available, where one is for free, the first extension at cost 1 adds two ports, and the last port is obtained by a second extension at a total cost of 3. Moreover, not more than three channels can be used on the original link between nodes 1 and 2. There is one spare channel for free, while the first extension at cost 1 adds the other two channels. The two additional small arcs at each node simply indicate the direction of all other adjacent arcs and does not have to do with the represented capacities.

For the considered demand, the rerouting of the lightpaths is now derived by the the solution of a minimum cost flow problem in the constructed network. The number of lightpaths corresponds to the total flow which begins in the target copy of the original source node of the demand and ends in the source copy of the original demand target node. Since all capacities are integer, the resulting flow is integer as well (see [1]) and hence represents a desired rerouting of the former fractional routing part for the demand at hand. It is to mention that in general, the flow may contain cycles with zero total cost. These cycles are removed in order to avoid unnecessary detours for the lightpaths. Note that the deletion of those cycles does not affect the feasibility of the remaining flow. Since the flow has minimum cost and all arc costs are nonnegative, no other cycles occur. Finally, we can use any path decomposition of the flow to identify the routing of the single lightpaths.

Because the complete verification proof for the adapted procedure is beyond the scope of this report, we restrict on some remarks for this. First of all, the diversification parameter setting guarantees for the existence of a solution for the minimum cost flow problem. By construction, the capacities modeled by the arcs ensure that the resulting total demand routing satisfies the diversification conditions for all nodes and links. Therefore, the procedure results in a feasible integer lightpath routing. In addition, note that each original link is only used in one direction by the rerouted lightpaths since the computed flow does not contain cycles. Moreover, it should be noted that the cost of a min cost flow do not correspond with the actual increase of the network cost by the rerouting. Without increasing the complexity of the rerouting problem substantially (i.e., to an NP-complete problem) the expansion cost can not be modeled correctly in general. The problem is twofold here. First, the actual expansion cost are fixed charge, i.e., the cost does not increase linearly with the used capacity but is a step function. Such a function can not be represented within a minimum cost flow problem. Second, in case of multiple expansion steps it can be the case that the incremental cost is lower for a higher capacity level (e.g., first expansion costs 10, second expansion only 5). In case only the incremental cost is set as cost coefficient in the minimum cost flow problem, the cheaper expansions would be used first. Therefore, we decided to set the coefficients to the total increase of the cost. In this way, the expansion levels are used in the right order.

Finally, we point out that the proposed cost model represents the actual expansion cost correctly in case that every node/edge needs a capacity expansion for at most one lightpath. Hence, in case only one lightpath has to be rerouted for a commodity, the cost of the min cost flow are the actual expansion cost. In fact, in this case the previously used shortest path approach and the min cost flow approach are equivalent. Note also, that both approaches have polynomial running time and that the min cost flow approach is better than the shortest path approach in the sense that multiple lightpaths (of the same commodity) are rerouted simultaneously. As before, the order in which the demands are rerouted can influence the total additional cost. However, since this cost is small in comparison with the total dimensioning cost, the implemented method works satisfactory, but can be refined if necessary.

2.2 Wavelength assignment

The continued research on the wavelength assignment subproblem focused on the two approaches which have been presented in the last report [6, Section 2]. On the one hand, we have carried out further investigations for the exact solution approach. On the other hand, the improvement heuristic has been reconsidered to exploit different variations.

2.2.1 Exact solution approach

The exact solution approach for the minimum converter wavelength assignment problem bases on an alternative formulation which has been introduced in [5, Section 3.2]. This formulation avoids the spectral symmetry involved by the commutability of the wavelengths. Such a symmetry causes a high degeneracy of a problem formulation which distinguishes between assignments differing only from a permutation of the wavelengths. As effect, the linear relaxation of the associated integer program includes fractional combinations of such (possibly even infeasible) assignments. These combinations often yield better solutions of the linear relaxation, but are infeasible for the underlying integer program and hence undesirable. To overcome this effect, the alternative formulation relies on a model in which only different assignments are explicitely distinguished. The presented integer program has a linear relaxation which excludes the unwanted combinations and hence results in better lower bounds.

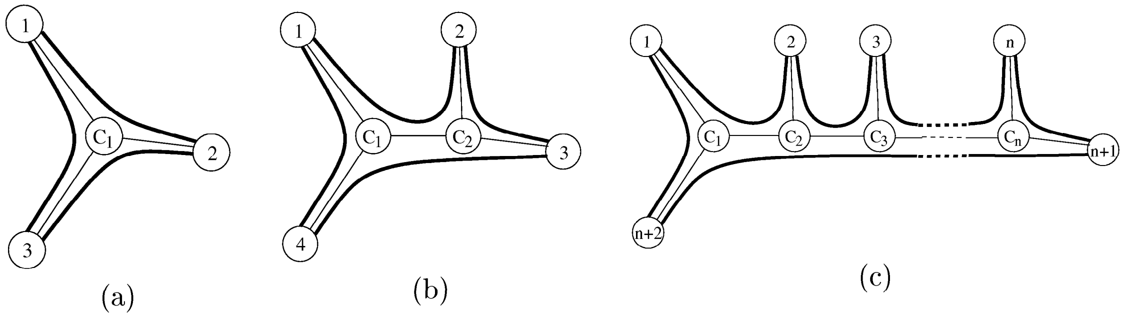

(a) (b) (c)

Figure 3: Instance series for the minimum converter wavelength assignment problem.

In fact, it has turned out that the lower bound improvement by use of the alternative formulation can be arbitrarily large compared to the standard formulation presented in [4, Section 3.3.2]. For proving this, we consider a series of instances for the wavelength assignment problem as illustrated in Figure 3. The thin lines denote the network links, while the thick lines represent the lightpaths. Each link provides the two wavelengths λ_1, λ_2 once. We begin with the left instance (a). Since both lightpaths emanating node 1 must use different wavelengths on the first link and are conflicting with the lightpath between node 2 and node 3, either one of them or the latter needs to change the used wavelength in the center node c_1. It is easy to see that any optimal wavelength assignment requires to place one converter in node c_1. Without going into detail, such an assignment is also indicated by the solution of the linear relaxation of the alternative formulation which thus provides the optimum value 1 as lower bound. For the standard formulation, the linear relaxation yields the optimum value 0 by assigning each wavelength half to each lightpath and hence avoids the use of converters. Note that this solution is a fractional combination of the two (infeasible) assignments in which all lightpaths are assigned the same wavelength.

Turning to instance (b), a similar conclusion as before shows that any optimal wavelength assignment requires two converters to be placed in the nodes c_1 and c_2. In this case, the linear relaxation for the alternative formulation indicates 2 as lower bound, while the for the standard formulation again yields 0 for this by assigning each wavelength half to each lightpath. Now, an induction argument can be applied. We extend the network stepwise as illustrated in the right picture (c) in Figure 3. For any $n \in \mathbb{Z}^+$, we get n and 0 as lower bounds from the linear relaxation of the alternative and the standard formulation, respectively, and that concludes the proof.

Although this investigation may seem to be mainly of theoretical interest, there are two important implications. First, the derived construction may give valuable hints for the identification of (partial) lightpath routing structures for which a nonzero lower bound can be simply deduced. We remark that the lightpath conflict graph of the general instance (c) in Figure 3 is very similar to a so-called *n-wheel* which plays an important role for the vertex coloring problem in graphs. Such an approach is especially recommended for large problem instances which are not tractable by integer programming techniques, and thus provides an interesting direction for further research. And second, if integer programming is used, the result strongly supports the application of the alternative formulation from a mathematical point of view.

Consequently, we have continued to investigate this approach. In the previous milestone, we have already reported on preliminary studies for the solution of the alternative integer program by a column generation method. A first implementation has turned out that the method succeeded in solving the linear relaxation, but left two questions open. On the one hand, the convergence speed of the column generation was quite slow for large problem instances. On the other hand, the solution of the linear relaxation is often characterized by many fractional variable values and hence has to be processed further on. Both aspects have been investigated during the ongoing work.

For improving the performance of the column generation method, several tuning possibilities have been examined. A main progress was achieved by an advanced initialization of the program. Instead of beginning with a single column which only contains all single links as partial paths and hence yields the worst possible solution in the first step, we can use the best (or any) wavelength assignment that has been derived by our heuristics. This assignment is transformed into an appropriate set of columns which are initially inserted into the linear program. As result, the column generation procedure starts at a far better point and hence reaches the final linear solution much faster.

A further speed-up was motivated by the obervation that the solution bases on a small subset of all partial paths in most cases. To take advantage of this insight, we begin with a reduced program which includes only those partial paths which are most probably needed. The excluded partial paths are then also priced out in each step and inserted if required. This additional effort mostly pays off by the reduction of the linear program which has to be processed in each step. Another idea concerns the objective value decrease which is achieved by each insertion of a generated column. It has been observed that this decrease becomes smaller and smaller during the processing, leading to only tiny objective improvements towards the end. We tried out to enhance the progress by several pricing step variations which aim at the generation of multiple improving columns instead of just one. This approach intends to reduce the total number of pricing steps and to enlarge the chance of higher value decreases. Computational experiments have shown this to work well at the beginning, but at the cost of a much faster growing size of the linear program which therefore needs more and more time to be resolved. In addition, the insertion of many similar columns can result in a bad quality of the derived solution, expressed in terms of the number of fractional column variables. (This effect bases on the existence of many optimal solutions using different column sets to cover the needed partial paths.)

To deal with fractional solutions of the linear relaxation, advanced techniques are required. A common approach is known as *branch-and-bound*. For a detailed description of the general method, we refer to [8]. The application of branch-and-bound requires to find a good branching rule which serves to split the problem at hand into suitable partial problems. For this, we have explored several possibilities. A simple and often used rule consists in branching on a fractional variable whose value range is constricted in each subproblem such that the current fractional value can not be taken anymore. This rule performs best for programs with a small number of fractional variables, i.e., a current solution which is already almost integer. Since the linear relaxation of our program is often characterized by many fractional variables, it is not recommended to apply this branching. Mehrotra and Trick [7] developed an efficient rule in their similar approach for the vertex coloring problem. Their rule unfortunately does not carry over to our problem.

Recently, Fischetti and Lodi [3] have published a new basic idea called *local branching* which is promising for the minimum converter wavelength assignment problem. Local branching intends to quickly find good integer solutions and to improve them as fast as possible. For this, the considered problem is split into two parts: a small neighborhood of a reference solution and the remaining problem. The smaller problem part has to be chosen in such a way that the associated integer program can be solved efficiently. At a first stage, arbitrary reference solutions can be used until a first feasible integer solution is at hand. Further on, the (best) integer solutions found so far serve as reference since it is likely to find better solutions nearby. This method aims at a quick improvement of the upper bound with which hopefully large parts of the remaining problem can be cut off.

A main advantage of local branching for the minimum converter wavelength assignment problem consists in the applied partition of the problem. By solving the smaller part immediately after generation, only the remaining program has to be maintained throughout the processing, reducing the computational effort especially in conjunction with column generation. Moreover, a fast upper bound improvement is best suited to support the solving of the remaining problem. Even if this hard part can not be solved to optimality, the approach is at least likely to provide good solutions.

For the application of local branching, it is necessary to define neighborhoods of wavelength assignments which are on the one hand small enough to allow for an efficient processing and on the other hand sufficiently large to probably contain advanced solutions. We have tested different possibilities for such a definition. As an example, the use of at least $|\Lambda| - k$ columns from the reference solution can be forced by an additional constraint. In the associated neighborhood, all (integer) solutions differ by at most k wavelengths from the reference assignment. The size of such a neighborhood grows considerably with the parameter $k \in \mathbb{Z}^+$, and it is not clear so far how to adjust this value in the best way. The appropriate search again requires to take all possible (not yet included) columns into account and has therefore also to be solved by a column generation procedure. First computational experiments have shown that this step often either fastly yields the unsatisfactory result that the neighborhood does not contain a better assignment, or takes long time to derive the optimal solution (of this smaller partial problem). Similar observations have been made for alternative definitions. Hence, further research is necessary to derive suitable neighborhoods. For the remaining problem part, an additional direction for advancements is the improvement of the lower bound which serves as benchmark for the quality guarantee for the best found wavelength assignment and is finally required to prove that the optimum has been reached. Unfortunately, it is not easy to construct moderate sized sample instances which are nontrivial and have a manageable number of lightpaths as well as a nonzero optimum. Investigations based on such instances are highly expected to provide more insight into the problem. The so far performed experiments indicate that this exact solution method is a promising alternative to the heuristic methods.

2.2.2 Improvement heuristic

Complexity considerations documented in the last milestone report have shown that the minimum converter wavelength assignment problem is hard to solve to optimality and thus recommend the development of good heuristics. In addition to the previously presented greedy methods, we also derived and discussed an alternative heuristic approach based on iterative improvements.

First computational tests indicated that this approach is very promising. Consequently, this investigations have been continued during the last project phase. The iterative improvement approach represents a general method which allows for multiple variants. In order to increase the performance as good as possible, we have generated and implemented several of these variants and compared their results on a series of wavelength assignment instances.

Basically, the iterative improvement method processes all lightpaths sequentially in each iteration. Each lightpath gets the wavelengths assigned in the best possible way, i.e., beginning with the first link, we consecutively choose a farest reaching wavelength. After the sequence is finished, the lightpaths are reordered, and the next iteration begins. The method terminates if a given time limit is exceeded.

The way the sequence is processed offers a first option for improvement. In the static version, each lightpath is processed immediately, independent of the result of the assignment. The alternative idea tries to avoid the use of converters as long as possible and not to block wavelengths by assigning them to lightpaths which anyway need converters. For this, the dynamic version postpones these lightpaths and processes them at the very end.

At second, the lightpath sequence can be reordered in different ways. On the one hand, all lightpaths with wavelength converters can be pushed in front of the sequence. On the other hand, we can restrict on moving only a subset of those lightpaths, for instance the first or the last one.

A third difference is expressed by the tie breaking rule for the choice of the assigned wavelength if several are possible. A variety of such rules can be applied. A very simple method consists of taking the first available wavelength according to a fixed ordering of the spectrum. A more sophisticated way is to keep track of the formerly used wavelengths and try to find a better choice according to this information. For instance, it is reasonable to expect that the assignment of heavily used wavelengths leaves most flexibility for the remaining spectrum. Another idea recommends to take a wavelength which is the most (or least) number of times available on the (partial) path to be assigned to. Finally, we have also considered versions which backtrack the wavelength use on the path before the current link and choose a wavelength which reaches the minimum number of links. Clearly, the presented variants do by far not cover all possibilities.

Computational studies of these variants have shown that (on average) the best choice consists of applying the following options:

- static sequence processing;
- all lightpaths using wavelength converters are pushed in front of the sequence after each iteration;

instance	Old computations best greedy	 iterative	New computations best greedy	 iterative
n17-s04-2g-d0	0	0	0	0
n17-s04-2g-d100	86	4	88	4
n17-s04-10g-d0	0	0	0	0
n17-s04-10g-d100	34	14	20	14
n50-s04-2g-d0	96	15	31	0
n50-s04-2g-d100	501	193	273	19
n50-s04-10g-d0	42	31	17	7
n50-s04-10g-d100	61	1	88	8

Table 1: Comparison of the best greedy heuristic and the iterative improvement heuristic results.

- if multiple wavelengths can be used, choose the first which has been mostly used so far.

Since the concrete values are of minor interest, we ommit a detailed description of the performed tests. Instead, the improvement gained with respect to the greedy heuristics is documented in Table 1. First, we have applied the iterative improvement heuristic on the transparent scenario instances reported in the last milestone (indicated by 'old computations'). Since these computations have been meanwhile renewed on changed settings, we have carried out the same comparison on these instances, too (indicated by 'new computations'). For all tests, a time limit of 10 minutes for the iterative heuristic was set. The results show that throughout all experiments, the number of converters is substantially reduced by the iterative heuristic. Moreover, note that for one instance the installation of converters has become unnecessary.

3 Computational studies

As a main emphasis of the second project part, the computational studies have been intensified and extended. By use of the developed solution method, we have computed static optical network configurations for a variety of different settings. In this section, we describe the computational experiments performed during the last half year and discuss the gained results. First of all, we revisit the computational study reported in [6] and compare the result with a software tool tested by T-Systems Nova. Next, we present the results of a case study where long haul and ultra-long haul transmission systems are compared on their implications for the routing and wavelength assignment. Finally, we present a computational study on the consequences of different levels of protection in optical networks.

3.1 Comparison of different technological scenarios

First computational comparisons for network designs based on different optical technologies have been reported in the previous milestone [6, Section 4]. The goal of such studies is to gain insight into possible future evolution and trends of optical networks. In particular, it is of interest to exploit the consequences induced by technological decisions from an operators point of view.

In order to keep our settings as realistic as possible, we have revised the device specifications as well as the underlying cost model in close cooperation with our colleagues from T-Systems Nova. The applied changes and adaptions required to actualize the computations and the discussion of the results. In addition, we compare our results with those derived by another software tool for computing static optical network configurations.

3.1.1 Instances

A basic introduction of all considered instances can be found in [6, Section 4.1]. Instead of repeating most of the details, we restrict on the description of the applied changes. All characteristics and parameters carry over if not mentioned explicitely in the following.

The major change concerns the cost model for the devices. In the opaque scenarios (s01), the WDM systems have a base cost of 24 and are additionally equipped with transponders at both fiber ends for each used channel. Such a transponder pair has a cost of 1.2 for 2.5 Gb and of 4 for 10 Gb. Thus, each installed WDM system has an individual total cost depending on its load. A similar model applies to the OXCs for which we distinguish three types O1, O2, and O3. In addition to a base cost for each type, the used ports have to be equipped with interfaces, also leading to individual total costs according to the number of switched lightpaths. Moreover, the port numbers differ for the bitrates of 2.5 Gb and 10 Gb. Since the switching is performed electronically and the backplanes have a fixed total bitrate which can be handled, each port for a 10 Gb lightpath is realized by using 4 ports operated at 2.5 Gb. As a result, the OXC types provide 256, 512, and 1024 ports in the 2.5 Gb cases and 64, 128, and 256 ports in the 10 Gb instances for O1, O2, and O3,

	2.5 Gb bitrate			10 Gb bitrate		
	O1	O2	O3	O1	O2	O3
OXC port number	256	512	1024	64	128	256
OXC base cost	85.0	125.0	190.0	85.0	125.0	190.0
cost per interface	1.2	1.2	1.2	4.8	4.8	4.8
full equipped OXC cost	392.2	739.4	1418.8	392.2	739.4	1418.8
total cost for OXCs with 17 used ports	105.4	145.4	210.4	166.6	206.6	271.6
	2.5 Gb bitrate			10 Gb bitrate		
WDM system base cost	24			24		
WDM cost per channel	1.2			4		
full equipped WDM cost	72			184		
total cost for WDM system with 17 channels	44.4			92.0		

Table 2: Component costs and examples for the opaque scenarios (s01).

scenario	WDM cost	OXC cost			transponder cost (regenerator/ wavel. conv.)
		O1 256 ports	O2 512 ports	O3 1024 ports	
s02-2.5G	72	244.4	414.8	724.6	–
s02-10G	184	244.4	414.8	724.6	–
s04-2.5G	24	244.4	414.8	724.6	0.6
s04-10G	24	244.4	414.8	724.6	2.0

Table 3: Component costs for the semitransparent (s02) and transparent (s04) scenarios.

respectively. The interface for a 2.5 Gb port causes a cost of 1.2, and thus the 10 Gb port interface has a total cost of 4.8 each. Table 2 displays some characteristic cost values for the opaque scenario.

In contrast to the opaque case, the OXCs in the semitransparent (s02) and transparent (s04) instances can not be individually equipped with interfaces. Therefore, each single OXC device of each type has the same cost as listed in Table 3. The OXCs perform optical switching which is independent of the bitrate and thus provide 256, 512, and 1024 ports in all cases. In addition to the base costs of 142.0, 210.0, and 315.0, respectively, each port requires an optical interface at a cost of 0.4. For the WDM systems, the same as before holds for the semitransparent cases. Each used channel must be equipped with a pair of transponders at the same costs as above. Table 3 displays only the total cost of fully equipped WDM systems for the s02 instances. In the transparent scenario, transponders are not required at the WDM systems which therefore incur only the base cost, but may be placed for regeneration as well as for wavelength conversion at a cost of 0.6 for 2.5 Gb and of 2.0 for 10 Gb lightpaths.

instance	nodes	edges	data rate	survi- vable?	total demand
n17-2.5G-d0	17	26	2.5 Gb	No	2662
n17-2.5G-d100	17	26	2.5 Gb	Yes	5324
n17-10G-d0	17	26	10 Gb	No	686
n17-10G-d100	17	26	10 Gb	Yes	1372
n50-2.5G-d0	50	88	2.5 Gb	No	2365
n50-2.5G-d100	50	88	2.5 Gb	Yes	4730
n50-10G-d0	50	88	10 Gb	No	904
n50-10G-d100	50	88	10 Gb	Yes	1808

Table 4: Characteristics of the instances.

Beside the described adaptions, all other settings remain unchanged. Table 4 recalls some instance characteristics which are independent of the technological scenario.

instance	# used edges	total # WDMs	# OXCs			# transponder	total cost	opt. gap
			O1	O2	O3			
n17-s01-2.5G-d0	25	111	11	3	5	–	19782.01	–
n17-s01-2.5G-d100	26	308	5	11	12	–	51016.38	–
n17-s01-10G-d0	22	34	12	4	4	–	15903.64	–
n17-s01-10G-d100	26	81	2	10	14	–	37443.62	–
n17-s02-2.5G-d0	25	111	11	3	5	–	16789.41	–
n17-s02-2.5G-d100	26	308	5	11	12	–	40389.98	–
n17-s02-10G-d0	22	34	15	2	0	–	10539.24	–
n17-s02-10G-d100	26	81	14	5	0	–	20899.22	–
n17-s04-2.5G-d0	25	111	12	2	5	0	11525.78	7.2%
n17-s04-2.5G-d100	26	299	6	13	11	7	25650.72	30.9%
n17-s04-10G-d0	22	34	15	2	0	0	5703.24	5.3%
n17-s04-10G-d100	26	82	15	5	0	0	8675.05	12.2%
n50-s01-2.5G-d0	86	193	34	16	0	–	30460.92	–
n50-s01-2.5G-d100	88	466	30	29	5	–	68868.50	–
n50-s01-10G-d0	75	101	33	26	0	–	38137.66	–
n50-s01-10G-d100	88	214	24	30	20	–	81628.14	–
n50-s02-2.5G-d0	86	193	34	16	0	–	29165.32	–
n50-s02-2.5G-d100	88	466	30	29	5	–	57962.90	–
n50-s02-10G-d0	74	101	50	0	0	–	28064.09	–
n50-s02-10G-d100	88	216	36	14	0	–	50577.26	–
n50-s04-2.5G-d0	84	194	36	14	0	17	20339.18	12.7%
n50-s04-2.5G-d100	88	472	31	30	5	12	37594.06	31.4%
n50-s04-10G-d0	77	100	50	0	0	25	15184.88	5.5%
n50-s04-10G-d100	88	220	39	11	0	2	20602.48	18.5%

Table 5: Results of the optimization for all instances.

3.1.2 Results

We have recomputed all 24 instances which have been considered before. The applied solution method includes all advancements described in this report. The new results are presented in Table 5 which displays some characteristic values for the solutions.

The first column shows the number of used edges in the established topologies. Note that all potential links are used in each survivable network, while all non-survivable instances use less links. Moreover, the link numbers for the different scenarios of each instance are exactly equal in the small network and do not vary much in the large one. A similar observation holds for the number of installed WDM systems and OXCs of all types listed in the next four columns of the table. Except for the opaque scenario in which the port number of the OXCs has been reduced for the 10G bitrate, there are only marginal differences for the numbers of installed devices in comparable networks on different scenarios. Indeed, the configurations are very similar. This is mainly due to the fact that the fiber costs are relatively small in comparison to the WDM system costs. Hence, the length of links does not play an important role, and thus the costs of all links are almost equal. As a consequence, the objective in the different scenarios differs only by a fixed factor, which explains the similarity of the results.

Another observation concerns the integration of full survivability into the network. This requires to increase the total amount of capacities on average by a factor of slightly higher than two, an effect which has also been observed in the previous study. Moreover, Table 5 shows that the number of transponders in the transparent scenario play only a minor role. In comparison with the number of lightpaths that have been routed, only few transponders are necessary, and the incurred costs are neglectable.

In the last two columns, the total network costs and the relative difference to the lower bound are given. For the opaque and transparent scenario, the gap between the lower and upper bound can not be applied since only fully equipped devices can be taken into account in the optimization. Hence, the actual costs are computed afterwards based on the number of channels and ports used. In the transparent non-survivable instances, the gap is small which proves that the appropriate solution is (at least) close to the optimum. The gaps for the survivable cases are higher since a good estimation of unavoidable costs (as lower bound) is more difficult for a routing which has to be diversified.

For a subsuming comparison of different technologies, the results show a clear trend: throughout all instances, the network costs decrease with increasing scenario. Note that the cost savings are substantial for most of the cases. As an example, the n50-s01-2.5G-d100 network becomes 16% cheaper in the semitransparent scenario and 45% cheaper in the transparent scenario. For the 10G bitrate, the relative savings are even larger. In both networks the 10G-d100 instance yields the highest rate, a total cost difference of around 75% between the opaque and the transparent scenario. Hence, the results strongly recommend the installation of transparent optical networks.

3.1.3 Comparison

In order to evaluate our results, the same instances have been additionally processed by our colleagues of T-Systems Nova using the software tool 'WDM Guru'. Table 6 lists for both approaches the resulting costs and their decomposition in node and edge costs for all instances. The difference is expressed relative to the WDM Guru results, i.e., the total cost of the OND result minus the of WDM Guru is divided by the latter. Consequently, the relative difference gives the portion of the total cost calculated by WDM Guru which are additionally needed (if positive) or safed (if negative) by OND.

For the opaque and semi-transparent scenarios, Table 6 shows that our optimization tool OND reports slightly worse results than WDM Guru. For the transparent scenario the results are comparable or a little bit better. These results can be explained by having a more detailed look at how OND works. OND is designed as decision support system for transparent optical networks. In these networks, the cost of a configuration depends on the devices installed. These devices typically have a certain capacity for transmitting/switching optical channels and the cost only depends on this capacity. In the opaque and semi-transparent scenario, however, the cost of these devices not only depends on the capacity provided but also on the number of channels actually used. In fact, the part of the cost that depends on

instance	WDM Guru			OND			relative difference (total cost)
	edge costs	node costs	total cost	edge costs	node costs	total cost	
n17-s01-2.5G-d0	9165.22	10475.2	19640.42	9233.61	10548.4	19782.01	0.7%
n17-s01-2.5G-d100	24877.60	24307.8	49185.40	25909.98	25106.4	51016.38	3.7%
n17-s01-10G-d0	5634.69	9356.0	14990.69	6043.64	9860.0	15903.64	6.1%
n17-s01-10G-d100	15142.24	21649.0	36791.24	15403.62	22040.0	37443.62	1.8%
n17-s02-2.5G-d0	9165.22	7554.0	16719.22	9233.61	7555.8	16789.41	0.4%
n17-s02-2.5G-d100	24860.20	14610.0	39470.20	25909.98	14480.0	40389.98	2.3%
n17-s02-10G-d0	5634.69	4490.0	10124.69	6043.64	4495.6	10539.24	4.1%
n17-s02-10G-d100	15121.69	5623.0	20744.69	15403.62	5495.6	20899.22	0.7%
n17-s04-2.5G-d0	4144.42	7554.0	11698.42	4140.38	7385.4	11525.78	-1.5%
n17-s04-2.5G-d100	10933.60	14610.0	25543.60	10817.12	14833.6	25650.72	0.4%
n17-s04-10G-d0	1258.16	4490.0	5748.16	1207.64	4495.6	5703.24	-0.7%
n17-s04-10G-d100	3123.97	5623.0	8746.97	2935.05	5740.0	8675.05	-0.8%
n50-s01-2.5G-d0	13939.69	15806.8	29746.49	14218.92	16242.0	30460.92	2.4%
n50-s01-2.5G-d100	33507.36	32828.0	66335.36	34978.70	33889.8	68868.50	3.8%
n50-s01-10G-d0	14357.09	20437.0	34794.09	15734.66	22403.0	38137.66	9.6%
n50-s01-10G-d100	33398.79	43307.0	76705.79	35582.14	46046.0	81628.14	6.5%
n50-s02-2.5G-d0	13914.97	14936.0	28850.97	14218.92	14946.4	29165.32	1.1%
n50-s02-2.5G-d100	33504.67	24385.0	57889.67	34978.70	22984.2	57962.90	0.1%
n50-s02-10G-d0	14303.21	12200.0	26503.21	15844.09	12220.0	28064.09	5.9%
n50-s02-10G-d100	33322.98	13910.0	47232.98	35971.66	14605.6	50577.26	7.1%
n50-s04-2.5G-d0	5996.09	14594.0	20590.09	5723.38	14615.8	20339.18	-1.2%
n50-s04-2.5G-d100	13615.06	24214.0	37829.06	13941.66	23652.4	37594.06	-0.6%
n50-s04-10G-d0	3077.04	12200.0	15277.04	2914.88	12270.0	15184.88	-0.6%
n50-s04-10G-d100	6256.20	14081.0	20337.20	6504.08	14098.4	20602.48	1.3%

Table 6: Comparison of the results concerning total network costs.

this number dominates the total cost. Although it is possible to model these costs correctly within OND (with many different artificial devices), state-of-the-art optimization does not allow for this alternative representation, since the resulting optimization problem is computationally not tractable anymore. Therefore, like in the transparent case, we only consider fully equipped WDM systems in OND, and release the spare interfaces/transponders in the configuration in a post-processing step.

Moreover, the cost structure of scenario 's01' and 's02' is favorable for a tool like WDM Guru that routes lightpaths sequentially. In this way, the individual channel cost can be taken into account without any additional effort. As a consequence, this difference in the objective between OND and WDM Guru leads to the slightly worse results for OND for the first two scenarios.

For the transparent case, improvements of upto 1.5% are reported. Taking into account that the gap between lower and upper bound in this particular case is only 7.2%, this improvement is in fact 17% of the potential improvement based on the lower bound.

3.2 Long haul vs. ultra-long haul transmission

In all computations that have been carried out a length restriction of 1200 km has been applied, which corresponds with today's so-called long haul systems. Nowadays, also so-called ultra-long haul systems are available. This systems have the advantage that optical signals can be transported on far longer ways without regeneration. Disadvantages of these systems are however that only a part of the wavelength spectrum can be used for these systems and that the cost of installation are substantially higher than those for long haul systems.

In this section, we report on a computational comparison of both systems. The goal of this study is to exploit the potential provided by very far reaching optical transmissions in terms of saved (regeneration) components and its influence on the total network costs. In particular, it is to examine if potentially saved regeneration costs are suspended by additional wavelength conversion costs caused by blocked wavelengths.

For this study, we consider a new instance with 65 nodes and 84 (potential) edges. Since not all pairs of vertices are 2-connected, we only consider protection for these demands where at least two node-independent paths exist. Since the length restriction only plays a role in the (fully) transparent scenario, we restrict ourselves to this case. Moreover, only 10 Gb bitrate is considered. Similar to the previous studies, there is one type of fiber, one type of WDM system, one type of regenerator, one type of wavelength converter, and three types of OXCs providing 256, 512, and 1024 bidirectional ports. Most of the components have the same cost in both instances. Each fiber costs 0.01 per km with an additional cost of 6 for amplification after each 70 km range. The OXCs are fully equipped with transparent interfaces for 0.4 cost at both sides of each port, resulting in a total cost of 244.4, 414.8, and 724.6 for the sizes 256, 512, and 1024, respectively. Finally, a combination of a multiplexer and a demultiplexer unit at cost of 12 is required at each fiber end, so each WDM system has a total cost of 24. However, the WDM system provides different wavelength numbers depending on the transmission system: while 40 wavelengths are available at each WDM system in the long haul case, the WDM spectrum contains only 20 wavelengths in the ultra-long haul instance. The main difference between the systems is the maximum optical transmission length which restricts lightpaths to a length of at most 1000 km in the long haul and at most 4000 km in the ultra-long haul case. Despite this, all other assumptions made in previous studies carry over. Wavelength converters are always able to convert each wavelength into each other provided by the used WDM systems, and the specific wavelength does not play a role for regenerating the signal or for switching it within the OXCs. In addition, we assume that signal regeneration as well as wavelength conversion is currently performed by the same device, a so-called transponder, so the transponder cost reflects the cost for both components. Note that the transponder cost is different in both instances. Table 7 gives an overview of some instance-specific information.

An overview of the results is given by the following two tables. While Table 8 describes the results in terms of components used or installed in the computed networks, the resulting costs are displayed in Table 9. We remark that 65 regenerators in the unprotected and 96 regenerators in the protected long haul instance are unavoidable since the minimum distance for these demands already exceeds 1000 km. Note that in the protected case, the need for node-disjoint paths for diversified demands has an additional impact on the routing. In

instance	total demand	# wavelengths per WDM	max. transm. length	transponder cost
n65-s04-10G-d0-LH	1228	40	1000	2
n65-s04-10G-d0-ULH	1228	20	4000	3
n65-s04-10G-d100-LH	1945	40	1000	2
n65-s04-10G-d100-ULH	1945	20	4000	3

Table 7: Characteristics of the long haul and ultra-long haul instances.

instance	used edges	total WDMs	OXCs			rege-ner.	wav. conv.	light-paths
			256	512	1024			
n65-s04-10G-d0-LH	83	149	63	3	0	167	0	1395
n65-s04-10G-d0-ULH	82	252	63	3	0	0	19	1228
n65-s04-10G-d100-LH	84	271	56	10	1	689	0	2634
n65-s04-10G-d100-ULH	84	489	56	10	1	0	2	1945

Table 8: Optimization results for the long haul vs. ultra-long haul system comparison concerning numbers of components.

particular, the network structure forces to use very long alternative paths for this.

The results show that the use of ultra-long haul systems pays off in terms of the number of transponders required. While all regenerators are eliminated by this, the number of wavelength conversions remains quite small. This yields a decrease in total node costs. However, the total costs for ultra-long haul networks are higher due to the halved capacity of the WDM systems at same price. This results in more systems that need to be installed and thus in higher total edge costs.

In order to compare the computed networks with the same WDM system capacity, we fixed the routings and transformed the WDM systems to the size of the other instance, i.e., the LH instances have to use WDM systems with 20 wavelengths, while the ULH instances are changed to WDM systems with 40 wavelengths. To indicate this change, the instance name gets an appropriate suffix (e.g., LH20 means the LH solution with halved WDM system sizes, and ULH40 means the ULH solution using 40 wavelength per WDM system). After

instance	total cost	node cost	edge cost	lower bound	gap
n65-s04-10G-d0-LH	22141.68	16975.6	5166.08	20451.9	8.26%
n65-s04-10G-d0-ULH	25648.79	16698.6	8950.19	24324.8	5.44%
n65-s04-10G-d100-LH	29740.60	19937.0	9803.60	22987.2	29.38%
n65-s04-10G-d100-ULH	36450.37	18565.0	17885.37	29815.7	22.25%

Table 9: Optimization results for the long haul vs. ultra-long haul system comparison concerning costs.

reconfiguring the edges, the wavelength assignment was computed again. Due to the fixed routing, only the edge configurations and the number of wavelength converters changed. The results are displayed in Table 10. Note that the LH20 instances have to be compared with the original ULH instances, while the ULH40 instances have to be compared with the original LH instances. To summarize this comparison, all ultra-long haul networks have a lower total cost than the appropriate long haul networks.

instance	WDM # wav.	total # WDMs	# wav. conv.	edge cost	node cost	total cost
n65-s04-10G-d0-LH20	20	274	0	9613.18	16975.6	26588.78
n65-s04-10G-d0-ULH40	40	149	0	5253.57	16641.6	21895.17
n65-s04-10G-d100-LH20	20	504	0	18486.76	19937.0	38423.76
n65-s04-10G-d100-ULH40	40	265	0	9630.71	18559.0	28189.71

Table 10: Optimization results for the long haul vs. ultra-long haul system comparison for exchanged WDM system sizes concerning costs.

3.3 Different survivability levels

In [6], we presented a survivability concept for optical networks. Basically, the operator has to specify for each demand the number of lightpaths that should survive in case of a failure. By diversification it is then guaranteed that at least these lightpaths are indeed available in failure states. If L denotes the demand in normal operating state and S the number of lightpaths that have to survive, the number of lightpaths to be routed equals $L' = \max\{L, 2S\}$, which are routed with diversification parameter $D = (L' - S)/L'$.

So far, we reported only results for $S = 0$ (d0-instances) and $S = L$ (d100-instances). In this section, we study the impact of other values for S. The comparison has been carried out on the transparent 's04' instances described earlier. Except for $S = 0$ and $S = L$, we computed network configurations for $S = \lceil 0.25L \rceil$, $S = \lceil 0.5L \rceil$, and $S = \lceil 0.75L \rceil$, which are marked by the suffixes 'd25', 'd50', and 'd75', respectively. Table 11 shows the total demand, the total network cost, the cost increase relative to the total cost for the appropriate network without survivability, the lower bound, the gap between the total cost and the lower bound, and the number of converters needed in the solution.

The results in Table 11 indicate for the first three networks that the cost increase upto a survivability percentage of 50% is rather marginal in comparison to the increase from 50% to 100%. This can be explained by the fact that almost no additional spare lightpaths (none for even-valued demands, only one for each odd-valued demand) have to be installed for a survivability degree of 25% or 50%, as shown by the total demand sums. Hence, the included diversification requirements yield the main impact on the network configuration. It turns out that such a diversified routing can be realized by a small increase of the total amount of installed capacities.

In the last network, n50-s04-10G, there are many demands of size $L = 1$. Even for the lowest survivability degree of 25%, this implies $S = 1$ as well and thus needs to insert a larger

instance	total demand	total cost	increase (in %)	lower bound	gap (in %)	total # converter
n17-s04-2.5G-d0	2662	11525.8	-	10761.2	6.66	0
n17-s04-2.5G-d25	2667	12772.2	10.8	10952.5	14.25	0
n17-s04-2.5G-d50	2688	14297.5	24.0	11820.1	17.33	11
n17-s04-2.5G-d75	4038	20669.4	79.3	15481.9	25.10	6
n17-s04-2.5G-d100	5324	25650.7	122.6	19602.2	23.58	7
n17-s04-10G-d0	668	5703.2	-	5414.7	5.06	0
n17-s04-10G-25	699	6094.9	6.8	5539.8	9.11	0
n17-s04-10G-50	720	6452.3	13.1	5653.3	12.38	19
n17-s04-10G-75	1066	7641.7	34.0	6851.8	10.34	10
n17-s04-10G-100	1372	8675.1	52.1	7731.3	10.88	0
n50-s04-2.5G-d0	2365	20339.2	-	17763.6	12.66	17
n50-s04-2.5G-d25	2365	23160.6	13.9	18180.3	21.50	0
n50-s04-2.5G-d50	2452	23866.5	17.3	18877.6	20.90	4
n50-s04-2.5G-d75	4194	34797.3	71.1	24234.1	30.36	16
n50-s04-2.5G-d100	4730	37594.1	84.8	25784.3	31.41	12
n50-s04-10G-d0	904	15184.9	-	14352.8	5.48	23
n50-s04-10G-d25	1490	18569.3	22.3	16276.6	12.35	0
n50-s04-10G-d50	1521	18947.7	24.8	16333.0	13.80	13
n50-s04-10G-d75	1722	20067.3	32.2	16695.3	16.80	3
n50-s04-10G-d100	1808	20602.5	35.7	16791.8	18.50	0

Table 11: Cost comparison for different levels of survivability.

portion of spare lightpaths. As a consequence, the cost increase for the first survivability level is higher than for the other networks, but does not grow that fast for the next levels. Such a demand structure yields high 'setup' costs for involving any survivability, but then allows to improve the degree at lower additional costs. Finally, it can be observed that the cost increase for the high survivability levels in both 10G networks is substantially lower than for the 2.5G networks. This corresponds to the sizes of the total demands. In particular, we see that the total node costs increase marginally from 0% to 100% (see Table 6), which indicates that many of the OXCs are not fully occupied in the non-survivable case.

To subsume this computational study of different survivability levels for optical networks, we can draw the following conclusions. On the one hand, the demand structure has a main influence on the relation between the increase of the survivability degree and the appropriate cost increase. On the other hand, the developed survivability concept allows for a good trade-off between the achieved network reliability and the therefore required additional investments.

4 Conclusions

The goal of the work package "Optimization of the static configuration of optical transport networks" was to investigate optical networks with different technologies and survivability requirements. The purpose of the investigation focused on cost-based comparisons of the static network configurations in order to evaluate the differences from a network operators point of view.

As first step, a general framework for each considered network technology has been specified. Three scenarios have been worked out: opaque, semitransparent, and transparent optical networks. While opaque and semitransparent systems yield point-to-point WDM networks which only differ in the used switching devices, the transparent scenario results in 2nd generation optical networks which are characterized by lightpaths as ongoing optical connections. Our considerations mainly oriented on the transparent case which reflects the most recent developments in optics and requires to take additional planning aspects into account. For this most general scenario, a mathematical programming formulation has been derived. Based on this formulation, a solution method for minimizing the total cost of such static network configurations has been developed and implemented. As a result, the computation of optical networks for all scenarios made it possible to compare the technologies and their impact on the configuration costs. It turned out that transparent optical networks are most favorable.

Besides several enhancements of the solution method and ongoing comparisons of network configurations for changing technological specifications, the second project part was mainly dedicated to further computational studies and in particular to the integration of reliability aspects into the network configuration. Therefore, we investigated known survivability schemes and developed a new concept which is in particular suited for optical networks and provides scalable security. This concept has also been adapted in our solution method with which the impact of a varying survivability level for the configurations and costs has been examined. Moreover, a comparison of long haul and ultra-long haul transmission systems has been carried out.

Although the developed solution method has proven to be sufficiently flexible to cover many different optical network specifications and can be used to answer various questions, issues for further research can be pointd out. First investigations have shown the exact solution approach for the minimum converter wavelength assignment problem to be promising. Hence, it is advised to continue the development of this approach which is of particular importance for transparent optical networks. Furthermore, alternative survivability concepts can be evaluated. Another topic is to investigate possible improvements for the computed network configurations as well as the lower bound on their costs by integration of so far subsequently treated aspects within the optimization.

References

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications.* Prentice Hall, New Jersey, 1993.

[2] G. Dahl and M. Stoer. A cutting plane algorithm for multicommodity survivable network design problems. *INFORMS Journal on Computing*, 10(1):1–11, 1998.

[3] M. Fischetti and A. Lodi. Local branching. Technical Report OR-02-6, May 2002.

[4] A. Koster and A. Zymolka. Routing in optical transport networks – optimization of the static configuration of optical transport netowrks – report milestone 1. Technical report, Konrad-Zuse-Zentrum für Informationstechnik Berlin, March 2001.

[5] A. Koster and A. Zymolka. Routing in optical transport networks – optimization of the static configuration of optical transport netowrks – report milestone 2. Technical report, Konrad-Zuse-Zentrum für Informationstechnik Berlin, September 2001.

[6] A. Koster and A. Zymolka. Routing in optical transport networks – optimization of the static configuration of optical transport netowrks – report milestone 3. Technical report, Konrad-Zuse-Zentrum für Informationstechnik Berlin, March 2002.

[7] A. Mehrotra and M. A. Trick. A column generation approach for graph coloring. *INFORMS Journal on Computing*, 8:344–354, 1996.

[8] G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization.* Wiley, N.Y., 1988.

[9] R. Wessäly. *DImensioning Survivable Capacitated NETworks.* PhD thesis, Technische Universität Berlin, 2000.