Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ADRIAN ZYMOLKA
ARIE M.C.A. KOSTER
ROLAND WESSÄLY

Transparent optical network design with sparse wavelength conversion

ZIB-Report 02–34 (October 2002)

Transparent optical network design with sparse wavelength conversion*

Adrian Zymolka[†‡] Arie M.C.A. Koster[†] Roland Wessäly[†]

Abstract

We consider the design of transparent optical networks from a practical perspective. Network operators aim at satisfying the communication demands at minimum cost. Such an optimization involves three interdependent planning issues: the dimensioning of the physical topology, the routing of lightpaths, and the wavelength assignment. Further topics include the reliability of the configuration and sparse wavelength conversion for efficient use of the capacities.

In this paper, we investigate this extensive optical network design task. Using a flexible device-based model, we present an integer programming formulation that supports greenfield planning as well as expansion planning on top of an existing network. As solution method, we propose a suitable decomposition approach that separates the wavelength assignment from the dimensioning and routing. Our method in particular provides a lower bound on the total cost which allows to rate the solution quality. Computational experiments on realistic networks approve the solution approach to be appropriate.

1 Introduction

Today's optical networks have emerged in response to the fast increase of bandwidth demands in communication networks. Driven by the sophisticated needs of a multitude of users, more and more applications with growing data requirements force the operators to regularly upgrade their core networks. In this respect, optical communication technology outperforms traditional electronic transmission and switching in two major regards: speed and capacity.

As enhancement of point-to-point optical networks with wavelength division multiplexing (WDM), transparent optical networks perform optical switching which enables the realization of ongoing optical connections throughout several links without conversion to electronics. Such optical connections form lightpaths which require the assignment of a wavelength

*This work has been carried out in cooperation with and supported by T-Systems Nova GmbH and Telekom Austria AG.

[†]Konrad-Zuse-Zentrum für Informationstechnik Berlin, Takustraße 7, D-14195 Berlin, Germany.

[‡]Corresponding author. E-mail: zymolka@zib.de — URL: http://www.zib.de/zymolka/

Keywords: optical network design, wavelength conversion, integer programming

Mathematics Subject Classification (1991): 90B18, 90C10, 90C35, 90C90

to use. Within a lightpath, wavelength converters can be applied to exchange the wavelength which allows for a more efficient use of the installed transmission capacities. In practice, the adaption of such technological advancements within the configuration planning states a perpetual challenge for network operators. Hence, there is a need for the development of appropriate optimization tools for the design and expansion of transparent optical networks that meet all requirements at minimum cost.

So far, several aspects of optical network design have been studied (see for example [5], [6], [7], [2], and [9]). In most of these studies, routing and wavelength assignment play an important role, whereas the dimensioning of the network is disregarded or treated implicitly. An exception in this context is [3].

From a theoretical point of view, it is interesting to solve just the wavelength assignment problem. However, operators judge the quality of such an assignment by means of total network operations cost. To our knowledge, there is little to no research available which deals with the planning of minimum cost transparent optical networks. The models and solution approaches we present in this paper are among the first attempts to formulate *and* solve a mathematical model that integrates a series of planning steps which have traditionally been solved sequentially.

The solutions provided by our approach suggest to the planner (1) a physical topology, (2) fiber decisions and (3) an installation of wavelength division multiplexers for all chosen links, (4) optical cross-connects for all nodes, (5) a lightpath routing for all communication demands that satisfies desired protection requirements, and (6) an assignment of wavelengths including the placement of converters. Our model enables the planner to perform expansion planning for an existing network as well as greenfield planning.

To solve such large-scale planning problems we propose in this paper a decomposition approach which separates the wavelength assignment problem from the dimensioning and routing part. The former problem is solved by means of combinatorial heuristics, whereas the latter one is solved using branch-and-cut type algorithms from our network dimensioning tool DISCNET.

Is such a decomposition contradictory to our claim of an integrated planning? Or more generally: how does one judge the overall quality of solutions provided by decomposing a problem into a sequence of subproblems? A serious answer to these questions must be based on information how much better an optimal solution to the overall problem could potentially be. This is exactly what our approach provides in addition to the proposed solutions: a lower bound on the minimum total network cost. For every problem instance, our algorithm computes at run-time such a lower bound which indicates the quality of our solutions.

The remainder of this paper is organized as follows. Section 2 contains a detailed description of the network architecture, the optical devices, the configuration of lightpaths, and the design problem under consideration. In Section 3, we introduce the integer programming formulation of the problem. The solution approach is described in Section 4. For the proposed method, computational results on realistic networks are reported and discussed in Section 5. Finally, concluding remarks and directions for further research are discussed in Section 6.

2 Problem description

In this section, we give a technical description of the optical network technology and the design problem as worked out in cooperation with our industrial partners. First, we introduce the underlying network architecture and the available optical devices with their functional specifications. These devices provide the required capacities to establish (optical) connections in the network. Next, we discuss how to configure such connections to satisfy given traffic demands. This also includes a survivability concept for optical networks. Finally, we state the optimization problem that assists network operators dealing with the design and upgrade of optical networks on a regular basis.

2.1 Optical network architecture

Basically, optical communication networks consist of a set of nodes that are connected by fiber links. We distinguish two layers of a network, representing the physical and the logical linking of nodes, respectively.

The physical layer contains the hardware to transmit optical signals between adjacent nodes and to process the data traffic within the nodes. An undirected graph $G = (V, E)$ represents the *physical topology* with all potential physical links. Note that although theoretically possible, G is usually not a complete graph, but mostly based on an existing network structure, possibly extended by new links.

On top of the physical topology, *lightpaths* are established as logical connections between pairs of nodes. A lightpath consists of a bidirectional ongoing optical channel that connects two nodes $s, t \in V$ via a path of one or more consecutive physical links and which is optically switched in the nodes of the path. Lightpaths are further characterized by the wavelength used on each link. Since a lightpath allows for a direct data exchange between s and t, we say that the nodes are logically linked. As a consequence, each set of established lightpaths corresponds to a set of logical links, which constitutes the *logical topology*. Clearly, each logical link may contain multiple lightpaths connecting the same pair of nodes. In principle, lightpaths can be operated on different bitrates (e.g., 2.5 Gbit/s, 10 Gbit/s). However, a mixed use is mostly excluded by the devices. Thus, we assume that all lightpaths work with the same bitrate.

In the scope of transparent optical network design, the logical topology is predefined by the given connection demands that have to be satisfied. A lightpath configuration can be realized by the installation of optical equipment at the physical topology.

2.2 Optical devices

To establish lightpaths, transmission capacity has to be provided at the links as well as switching capacity in the nodes. In general, for each of these tasks several device types are available that differ in their capabilities and the underlying cost function. The following description of the devices focuses on their main properties with respect to network planning. For detailed technical explanations, we refer to [6].

Link technology

Transmission capacity on physical links is provided by *WDM systems* that have to be installed on top of *optical fibers*. Each link may contain (multiple) parallel fibers. The fibers are operated bidirectionally, i.e., they carry signals in both directions. Each WDM system occupies one underlying fiber and multiplexes a number of optical channels using different wavelengths into the transmitted signal. The WDM systems are also operated bidirectional. In practice, several types of fibers and WDM systems are available.

Let $\mathcal{F}$ denote the set of different fiber types and $\mathcal{W}$ the set of different WDM system types. Due to technical restrictions, some combinations are impossible, i.e., not every WDM system can be installed at every fiber type. The set of WDM system types installable on fiber type $f \in \mathcal{F}$ is denoted by $\mathcal{W}_f \subseteq \mathcal{W}$. Each WDM system provides a set of optical channels operated with specific wavelengths. Let Λ be the set of all possible wavelengths and $\Lambda_w \subseteq \Lambda$ the partial spectrum maintained by WDM systems of type $w \in \mathcal{W}$. To simplify the notation, the parameter $\kappa_w^\lambda \in \{0, 1\}$ indicates whether wavelength $\lambda \in \Lambda$ is provided by WDM system type $w \in \mathcal{W}$ or not. Note that WDM systems may in particular offer different spectrum sizes $\kappa_w = |\Lambda_w|$, referred to as their *channel capacity*. Finally, each of these device types has a specific cost function. Since the fibers have different lengths and must perform amplification in regular distances, their cost generally depends on the link of installation. In contrast, WDM system costs usually vary only by the types. Without specifying these functions in detail, the cost for installing a fiber of type $f \in \mathcal{F}$ on link $e \in E$ is denoted by $c_f(e)$ and for WDM system type $w \in \mathcal{W}$ by c_w.

Node technology

In the nodes of the physical topology, equipment for switching of lightpaths has to be installed. In addition, devices to convert the wavelength of a lightpath may be placed. These operations are performed by different devices.

In order to establish switching capacity, *optical cross-connects* (OXCs) must be installed. Each OXC provides a number of ports able to switch lightpaths from one link to the next. Let $\mathcal{O}$ denote the set of distinct OXC types which differ in the available port number. An OXC of type $o \in \mathcal{O}$ has a switch matrix of $\kappa_o \times \kappa_o$ ports and thus can handle κ_o lightpaths. The physical nodes can be equipped with multiple OXCs whose switch capacities are additive, i.e., there is no loss of available ports by device interconnections. Typically, the cost for OXCs is independent of the installation node. Hence, we denote the cost for an OXC of type $o \in \mathcal{O}$ by c_o.

Lightpaths are usually operated at the same wavelength along the path from source to sink. At the cost of installing a so-called *wavelength converter*, the wavelength of a lightpath can be exchanged at an intermediate node. In this way, the transmission capacity at the links can be deployed more efficiently, since a lightpath can use another available wavelength if the desired one is blocked by another lightpath. Wavelength converters are placed between the OXC ports and the subsequent links and can switch any arbitrary wavelength of a single lightpath to any other wavelength, i.e., the accessible spectrum is not limited. Since no economics of scale can be expected for devices that convert multiple lightpaths simultaneously, we consider only one type of wavelength converter at cost c_C which converts a single lightpath. Currently, wavelength exchanges are realized by the coupling of a receiver and a transmitter which performs optic-electronic-optic conversion. It is expected that optical conversion will be available in the future.

2.3 Lightpath configuration and survivability

Let $\mathcal{Q}$ denote the set of all communication demands (also called commodities). Each demand $q \in \mathcal{Q}$ is specified by a pair $s_q, t_q \in V$ of nodes that have to be connected by $d_q \in \mathbb{Z}^+$ lightpaths. To establish a lightpath, a routing in the physical topology has to be determined. Multiple lightpaths connecting the same pair of nodes need not to follow a common path, i.e., the routing for such demands can be bifurcated. In addition, at each link of a lightpath we have to identify the wavelength to use. For two consecutive links, the wavelength can be changed if a wavelength converter for this lightpath is installed at the connecting node.

In the design of telecommunication networks, reliability of the network under the occurrence of failures plays an important role. Especially in the core of communication networks where optical technology is typically deployed, breakdowns cause the interruption of large amounts of traffic. In this context, several survivability concepts have been developed over the years, as for instance path restoration or 1+1 protection. The concept of diversification, introduced by [1], is particularly suitable as survivability concept in transparent optical networks and is therefore adapted in this paper.

Diversification aims at limiting the number of connections interrupted during an outage by using an advanced routing in normal operation. For each demand $q \in \mathcal{Q}$, we introduce a parameter $0 < \delta_q \leq 1$. The lightpaths of q have to be routed in such a way that at most $\lfloor \delta_q d_q \rfloor$ of them cross the same physical link or node. In case a link or node (except of s_q, t_q) fails, the bound guarantees at least $d_q - \lfloor \delta_q d_q \rfloor$ connections of q to survive. Note that the parameters have to be chosen carefully since the physical topology must provide sufficiently many node-disjoint paths between the endnodes for such a routing. For example, a diversification parameter $\delta_q = 0.5$ for a demand of $d_q = 5$ lightpaths implies that at least $\lceil d_q / \lfloor \delta_q d_q \rfloor \rceil = 3$ node-disjoint paths must be used.

2.4 Optimization problem

The transparent optical network design problem consists of finding a minimum cost network configuration that satisfies all communication demands. The total cost of a network configuration incorporates the cost of new fibers, WDM systems, OXCs, and wavelength converters (these costs may also include maintenance cost). The decisions to be taken in the optimization tool can be divided into three parts: dimensioning, routing, and wavelength assignment.

The *dimensioning* consists of the choice and placement of optical devices, possibly extending previously installed equipment. We assume that the preinstalled devices can be used for free, but are fixed in place. For each link $e \in E$, let $\pi_e^f, \pi_e^w \in \mathbb{Z}_0^+$ denote the preinstalled fibers of type $f \in \mathcal{F}$ and WDM systems of type $w \in \mathcal{W}$, respectively. Each node $v \in V$ already contains a number $\pi_v^o \in \mathbb{Z}_0^+$ of preinstalled OXCs of type $o \in \mathcal{O}$, and $\pi_v^C \in \mathbb{Z}_0^+$ wavelength converters are given. Moreover, there are upper bounds $\mu_e^f, \mu_e^w, \mu_v^o \in \mathbb{Z}_0^+ \cup \{\infty\}$ on the number of newly installable fibers, WDM systems and OXCs of each type, respectively. An infinite upper bound states unlimited usability. We assume that wavelength converters can be attached to all installed OXC ports.

The physical topology has to be dimensioned such that enough transmission and switching capacity is available for the *routing* of all demands. Furthermore, the lightpath routing has to satisfy the diversification restrictions. Finally, the *wavelength assignment* to every link of each lightpath has to be carried out in such a way that a wavelength is not used more than available at the link, and the number of wavelength conversions at a node does not exceed the available number of converters.

3 Model

In this section, we present an integer programming formulation of the transparent optical network design problem. We first introduce the decision variables, then we state the objective function as well as all constraints imposed by the problem description.

3.1 Decision variables

There are two classes of decision variables corresponding to the different tasks involved in transparent optical network planning. One part concerns the lightpath configuration with routing and wavelength assignment, and the other part models the placement of new optical devices as well as their combination.

The routing of lightpaths is modeled by flow variables, which requires the demands to be directed. We define s_q to be the source and t_q as target of the demand $q \in \mathcal{Q}$. Note that the direction can be chosen arbitrarily since the lightpaths represent bidirectional connections. In addition, we have to specify the wavelength used on each link of a lightpath. Therefore, we introduce for each demand $q \in \mathcal{Q}$ the routing variables $r_{uv}^{q\lambda} \in \mathbb{Z}_0^+$ representing the number of lightpaths that cross link $uv \in E$ from u to v on wavelength $\lambda \in \Lambda$. To obtain a linear formulation, we count by $z_v^{q\lambda} \in \mathbb{Z}_0^+$ the number of lightpaths that change the wavelength from λ to any other in node v, for each demand $q \in \mathcal{Q}$, all nodes $v \in V \setminus \{s_q, t_q\}$, and each wavelength $\lambda \in \Lambda$. To formulate the constraint, it is not important to which wavelength is converted to.

To provide the required capacities for a lightpath configuration, new devices can be installed in the physical topology. For each link $e \in E$, we introduce variables $x_e^f, x_e^w \in \mathbb{Z}_0^+$ representing the number of new fibers of type $f \in \mathcal{F}$ and the number of new WDM systems of type $w \in \mathcal{W}$, respectively. For each allowed combination, the assignment of underlying fibers to used WDM systems is expressed by the variables $a_e^{fw} \in \mathbb{Z}_0^+$ that state how many fibers of type $f \in \mathcal{F}$ carry WDM systems of type $w \in \mathcal{W}_f$ on link $e \in E$. With these variables, the number of times a wavelength is available on a link is determined. For each network node $v \in V$, the number of new OXCs of type $o \in \mathcal{O}$ is given by the variables $y_v^o \in \mathbb{Z}_0^+$. Finally, we introduce auxiliary variables $z_v \in \mathbb{Z}_0^+$ which indicate the total number of new wavelength converters in node $v \in V$.

3.2 Objective function and constraints

The objective in optical network design is to minimize the total network cost

$$\min \sum_{e \in E} \left(\sum_{f \in \mathcal{F}} c_f(e) x_e^f + \sum_{w \in \mathcal{W}} c_w x_e^w \right) + \sum_{v \in V} \left(c_C z_v + \sum_{o \in \mathcal{O}} c_o y_v^o \right) \tag{1}$$

incurred by the installation of new devices. The optimization problem is subject to the following constraints.

The routing of each demand $q \in \mathcal{Q}$ is represented as flow from the source node $s_q \in V$ to the target node $t_q \in V$ according to the flow conditions:

$$\sum_{uv \in E} \sum_{\lambda \in \Lambda} r_{uv}^{q\lambda} - \sum_{vu \in E} \sum_{\lambda \in \Lambda} r_{vu}^{q\lambda} = \begin{cases} -d_q & \text{, if } v = s_q, \\ d_q & \text{, if } v = t_q, \\ 0 & \text{, otherwise} \end{cases} \qquad \forall\, q \in \mathcal{Q}, v \in V \tag{2}$$

The portion of each demand $q \in \mathcal{Q}$ crossing a single link or node (except of source and target) is restricted by the diversification parameters:

$$\sum_{\lambda \in \Lambda} \left(r_{uv}^{q\lambda} + r_{vu}^{q\lambda} \right) \leq \lfloor \delta_q d_q \rfloor \qquad \forall\, q \in \mathcal{Q}, e = uv \in E \tag{3}$$

$$\sum_{u \in V} \sum_{\lambda \in \Lambda} r_{uv}^{q\lambda} \leq \lfloor \delta_q d_q \rfloor \qquad \forall\, q \in \mathcal{Q}, v \in V \setminus \{s_q, t_q\} \tag{4}$$

The necessary wavelength converters in a node v are separately counted for each wavelength $\lambda \in \Lambda$ as difference between incoming and outgoing channels:

$$\sum_{uv \in E} r_{uv}^{q\lambda} - \sum_{vu \in E} r_{vu}^{q\lambda} \leq z_v^{q\lambda} \qquad \forall\, q \in \mathcal{Q}, \lambda \in \Lambda, v \in V \setminus \{s_q, t_q\} \tag{5}$$

These numbers sum up to the total number of wavelength converters needed in each node $v \in V$:

$$\sum_{q \in \mathcal{Q}} \sum_{\lambda \in \Lambda} z_v^{q\lambda} \leq \pi_v^C + z_v \qquad \forall\, v \in V \tag{6}$$

The switching capacity in each node $v \in V$ is provided by the installed OXCs and has to cover all lightpaths that contain the node, i.e., reaching that node or beginning there:

$$\sum_{q \in \mathcal{Q}} \sum_{\lambda \in \Lambda} \sum_{uv \in E} r_{uv}^{q\lambda} + \sum_{\substack{q \in \mathcal{Q}: \\ s_q = v}} d_q \leq \sum_{o \in \mathcal{O}} \kappa_o \left(\pi_v^o + y_v^o \right) \qquad \forall\, v \in V \tag{7}$$

As transmission capacity for the total flow, each wavelength $\lambda \in \Lambda$ must be available sufficiently many times on each link $e \in E$ for all lightpaths in both directions:

$$\sum_{q \in \mathcal{Q}} \left(r_{uv}^{q\lambda} + r_{vu}^{q\lambda} \right) \leq \sum_{f \in \mathcal{F}} \sum_{w \in \mathcal{W}_f} \kappa_w^\lambda a_e^{fw} \qquad \forall\, e = uv \in E, \lambda \in \Lambda \tag{8}$$

The available wavelengths are determined by the installed transmission equipment. On each link $e \in E$, the number of usable fibers of each type $f \in \mathcal{F}$ is given by the preinstalled and the new fibers:

$$\sum_{w \in \mathcal{W}_f} a_e^{fw} \leq \pi_e^f + x_e^f \qquad \forall\, e \in E, f \in \mathcal{F} \tag{9}$$

The same holds for all types of WDM systems used:

$$\sum_{f \in \mathcal{F}} a_e^{fw} \leq \pi_e^w + x_e^w \qquad \forall\, e \in E, w \in \mathcal{W} \tag{10}$$

Finally, all variables are integer and restricted by upper bounds:

$$\mu_e^f \geq x_e^f \in \mathbb{Z}_0^+ \quad \forall\, e \in E, f \in \mathcal{F} \tag{11}$$

$$\mu_e^w \geq x_e^w \in \mathbb{Z}_0^+ \quad \forall\, e \in E, w \in \mathcal{W} \tag{12}$$

$$a_e^{fw} \in \mathbb{Z}_0^+ \quad \forall\, f \in \mathcal{F}, w \in \mathcal{W}, e \in E \tag{13}$$

$$\mu_v^o \geq y_v^o \in \mathbb{Z}_0^+ \quad \forall\, v \in V, o \in \mathcal{O} \tag{14}$$

$$z_v^{q\lambda} \in \mathbb{Z}_0^+ \quad \forall\, q \in \mathcal{Q}, \lambda \in \Lambda, v \in V \tag{15}$$

$$z_v \in \mathbb{Z}_0^+ \quad \forall\, v \in V \tag{16}$$

$$r_{uv}^{q\lambda} \in \mathbb{Z}_0^+ \quad \forall\, q \in \mathcal{Q}, \lambda \in \Lambda, e = uv \in E \tag{17}$$

4 The solution approach

In this section, we provide a solution approach for the considered transparent optical network design problem, based on a decomposition into two smaller optimization problems which are solved subsequently.

4.1 The decomposition approach

The integer program presented in the preceding section has a practical drawback: even for moderate problem sizes, the large number of integer variables makes the program computationally intractable for state-of-the-art mixed integer programming solvers. This is mainly due to the inclusion of an integer multicommodity flow problem which models the lightpath routing as well as a coloring problem for the wavelength assignment. It is well-known that both are hard to solve exactly. Especially the high symmetry in the wavelength assignment part causes difficulties in solving the problem to optimality.

To relieve the computational effort without losing too much solution quality, we decompose the transparent optical network design problem into a dimensioning and routing subproblem on the one hand and the wavelength assignment subproblem on the other hand. In this way, the coloring problem is separated from the integer multicommodity flow part. Sparse wavelength conversion allows for this decomposition, since it guarantees a feasible wavelength assignment for every routing. Moreover, the dimensioning and routing subproblem has fundamental similarities to the network design problem studied for electronic and WDM point-to-point networks, and hence can profit from their results. In addition, theoretical results on the wavelength assignment subproblem give reason to expect that the required number of wavelength converters in multi-fiber networks will be limited (cf. [4]). In the following, we describe how we deal with these two subproblems.

4.2 Dimensioning and routing

In the dimensioning and routing subproblem, we determine the dimensioning of transmission and switching capacities together with the lightpath routing. Wavelengths are not distinguished here. Instead, we express the transmission capacity by the number of channels provided by the WDM systems. Hence, the outcome of this subproblem consists of an integer routing of the demands on a suitably equipped physical network. Note that as a consequence, the number of wavelength converters can not be determined at this stage.

The dimensioning and routing problem for telecommunication networks has achieved considerable attention in the last decade. A variety of solution methods and optimization tools has emerged from these studies. One such tool is DISCNET (see [8]) which was initially developed for the design of SDH networks. DISCNET exploits integer programming techniques to achieve two goals: to find good solutions and to derive a lower bound on the total cost of any network design that satisfies the stated requirements. This lower bound in particular provides a measure for the quality of the solutions.

To take advantage of these features, we transform the problem input into a suitable setting for DISCNET in a preprocessing phase. To this end, we generate the set of installable capacity levels for each physical link and node. The application of DISCNET then yields a dimensioning of the network according to these capacity levels together with a routing for the demands which satisfies the specified survivability constraints. Since the routing can be fractional, a postprocessing step is deployed in order to turn the routing into an integer one. We choose to resolve the dimensioning and routing problem in this way, since in practice the number of non-integer routings is small in comparison to the total number of demands as well as in comparison to the computational effort involved by requiring an integer routing within the optimization.

Preprocessing

In the preprocessing phase, all installable capacity levels have to be generated for each physical link and node with respect to the existing and installable devices. The levels form discrete sets of increasing capacities and costs. In case of unrestricted upper bounds on the device numbers, we use the total number of lightpaths to set a maximum capacity level.

Given a link, the possible combinations of fibers and WDM systems up to their maximum multiplicity have to be taken into account. Instead of enumerating all configurations, dynamic programming can be applied. Without going into detail, the basic principle consists of maintaining a list of all potential channel capacity levels each of which has an associated cost that is initially set to infinity. Clearly, zero channels are available for free. This list is then updated in a sequential procedure. In each step, a single fiber-WDM system pair is considered. We replace each capacity level cost by the minimum of two values: the previous cost, indicating that the current combination is not included into the configuration, and the cost of the appropriate lower capacity level plus the cost for the inserted pair. Thereby, we keep track of the already used and remaining free device numbers. In the end, the list contains the minimum cost for each capacity level. Figure 1 shows an example where the level costs are displayed by the dots. Note that some capacities may be dominated by higher levels which can be realized at lower cost. All dominated levels are finally deleted, and the resulting lower envelope marks the desired set of capacity levels with associated costs.

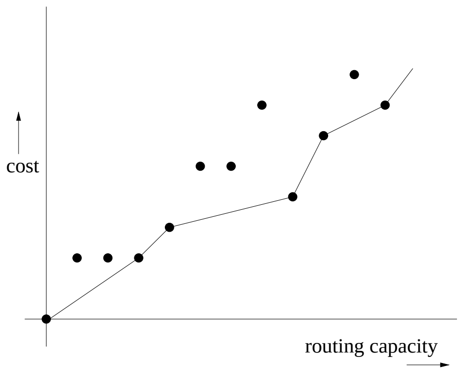

Figure 1: Example list of minimum cost for each capacity level.

The same method can be used for the switch capacity levels in the nodes. Instead of fiber-WDM system pairs, we proceed stepwise on all OXCs that can be installed in the considered node. Note that the final configuration for each capacity level can be acquired easily if we keep track of the inserted device(s) for each cost replacement. All installable capacities are generated by this method, i.e., the generated sets are describing all available options.

DISCNET

As a result, the computed sets of incremental capacity levels and costs can be composed to the input of DISCNET which uses an integer programming formulation similar to the one in Section 3. In a first stage, the linear relaxation is solved and a cutting plane algorithm is applied for several classes of valid inequalities. In this way, a lower bound on the minimum configuration cost is computed. The second stage intends to find good integer solutions using this strengthened linear relaxation. For this purpose, multiple branch-and-cut heuristics are deployed, which search only specific branches in depth instead of the whole branch-and-cut tree. Different branching rules result in different branch-and-cut heuristics. For more detailed information, we refer to [8].

Postprocessing

The output of DISCNET consists of the chosen capacity level for each physical link and node, the incurred total network cost, a lower bound on the cost for any feasible network configuration, and a routing of all demands which respects the diversification conditions. Since DISCNET does not require an integer routing, which is crucial for lightpath routing in optical networks, a correction procedure has to be applied in case the routing is indeed fractional. In this procedure, the fractional parts are rerouted.

The procedure starts by fixing the integer part of the routing. Note that the fractional parts of each demand sum up to an integer number of missing lightpaths. These integer demands are then processed sequentially. For each demand, an instance of the minimum cost flow problem is solved. The necessary capacities (on both nodes and links) for this flow are determined by the difference between the fixed part of the routing and the upper bound stated by diversification conditions. In addition, we distinguish between free capacity that is provided by so far unused channels, and additional capacity which is available by extension of the current capacity to the next level. Since the provided capacities are integer, we obtain an integer flow which represents the routing of the remaining lightpaths for the considered demand. If the flow cost equals zero, the lightpaths do not require the installation of additional transmission and switching capacity. Otherwise, we have to extend some installed capacities to the next level. This can also provide free capacities for the subsequent demands to process.

As a result of the postprocessing procedure, we obtain a network with sufficient transmission and switch capacity and an integer lightpath routing for which a wavelength assignment has to be computed to complete the lightpath configuration. We remark that the postprocessing can be extended to include further issues, e.g., a limited transmission length for lightpaths. In this case, so-called regenerators are installed to transform the configuration into a multi-hop lightpath routing.

4.3 Wavelength assignment

The solution of the dimensioning and routing subproblem yields the installed link configurations with appropriate channel capacities for the established lightpath routing. In order to assign wavelengths to the lightpaths, the channel capacities are reinterpreted as spectra of wavelengths provided by the used WDM systems. In general, such an assignment will require the placement of wavelength converters. Since the number of converters is not limited, given a routing, there always exists an assignment of wavelengths to the lightpaths. Hence, the task of this subproblem is to find a wavelength assignment with a minimum total number of wavelength converters. In this section we describe a number of algorithms to achieve this goal. For the sake of clarity, preinstalled wavelength converters are omitted in the description of these algorithms, but can be easily incorporated in each of them.

First of all, the wavelength assignment subproblem can be simplified by temporarily neglecting all lightpaths that use only a single link. After having assigned wavelengths to all other lightpaths, the established channel capacities guarantee that there will be sufficient wavelengths left for these lightpaths. Hence, all single link lightpaths are treated last.

Since each link used by a lightpath may have a different spectrum of available wavelengths, we first deal with the problem how to (optimally) assign wavelengths to a single lightpath. It is easy to see that the following method yields the optimal result. Beginning with the first link, we choose a wavelength that reaches as far as possible, i.e., which is available on the maximum number of consecutive links. If the whole lightpath is covered, we are done. Otherwise, we place a wavelength converter in the last node reached and repeat the procedure at the following link. Note that in any arbitrary assignment, the first wavelength converter cannot be placed later than above. Since this also holds for all next conversions, the above described method requires a minimum number of wavelength converters for a single lightpath.

A first heuristic for the wavelength assignment problem is to apply this method sequentially to all lightpaths. Of course, the order in which the lightpaths are processed has a major impact on the number of wavelengths needed. Moreover, for a single lightpath we have to break ties for choosing a wavelength if multiple are available. We applied several orderings, e.g., taking the longest lightpath first since it seems most complicate to be processed later on, or taking the lightpath with the fewest number of remaining wavelengths which are available on all links in order to decrease the danger of requiring an additional wavelength converter. For tie breaking, again a variety of rules can be applied, e.g., to take the first available wavelength according to a fixed ordering of all wavelengths, or to choose one that has been used at most links before. We tested several possibilities for both options and different combinations.

In addition to these greedy heuristics, we designed another heuristic which applies local search on the possible processing orders. If a lightpath requires a wavelength converter in the sequential procedure, it is probably processed 'too late'. Therefore, we push all such lightpaths to the beginning and repeat. In case a wavelength assignment without conversions is found, the heuristic terminates. Otherwise, we continue until reaching a time limit, and return the best assignment found so far.

5 Computational experiments

In this section, we report on computational experiments with the proposed solution approach. We first describe the test instances on which the study has been carried out. Then we present and discuss the results achieved with our method.

5.1 Instances

For all instances, we consider the same optical equipment. The underlying cost functions for the devices base on a realistic cost model whose values have been scaled by an appropriate factor. The parameters of the cost functions are displayed in Table 1.

We consider three different fiber types labeled by $\mathcal{F} = \{\mathrm{S,L,XL}\}$. The cost of each fiber is composed of a cost per km and an additional cost for required amplification after each segment of a specific length. The cost values and segment lengths vary with the fiber types. Which fiber type is cheapest depends on the distance. Moreover, two different WDM system

fiber type	km cost	seg. cost	seg. km	WDM type	chan-nels	chan. cost	fixed cost	OXC type	ports	port cost	fixed cost
S	0.1	0.6	70	W16	16	0.5	6	OXCS	128	0.6	50
L	0.1	1.0	100	W32	32	0.5	10	OXCM	256	0.6	90
XL	0.09	1.0	50					OXCL	512	0.7	120

Table 1: Cost function parameters for the optical devices.

types are available, denoted by $\mathcal{W} = \{\text{W16,W32}\}$. The system W16 can only be used in connection with fiber types S and L and provides 16 wavelengths, while the system W32 is installable on fiber types L and XL and provides the total spectrum of 32 wavelengths. The wavelengths available with W16 build a subset of those provided by W32. The cost of a WDM system is composed of a fixed base cost and a cost for each supplied channel (resulting in a total cost c_w of 14 and 26 for W16 and W32, respectively). A similar cost function is also applied for the OXCs. There are three OXC types in $\mathcal{O} = \{\text{OXCS,OXCM,OXCL}\}$ with 128, 256, and 512 ports, respectively. A fixed cost and a cost for each provided port sum up to the total cost of each OXC. Finally, each wavelength converter has a cost of 2.

We performed several computations on four realistic networks, partly provided by our industrial partners: the well-known NSF backbone with 14 nodes (nsf14), a 17 node German network (ge17), a 24 node network in the United Kingdom (uk24), and a European network with 42 nodes (eur42). All of them have a meshed physical topology. Some characteristic instance parameters are displayed in Table 2. Instead of limiting the maximum number for each single device type, we state an appropriate upper bound on the total number of fibers at each link and of OXCs in each node.

For each of these networks, we generated four problem instances which differ in the survivability concept and the preinstalled equipment. The suffix 'd0' indicates that diversification is not applied, i.e., no survivability is guaranteed. Instances marked with the suffix 'd50' require at least half of each demand to survive any link or node failure. Hence, the diversification parameter for each demand is set to 50% (all demands are even-valued). Note that the physical topology has to be 2-connected for this purpose. Moreover, we use the additional suffix 'pr' to point out the instances with preinstalled devices in which each link is equipped with one fiber of type S and one of type L as well as a WDM system of type W16. OXCs and wavelength converters are not provided. The instances without the suffix 'pr' do not contain any preinstalled devices, i.e., they represent optical network design on the greenfield.

instance	$\|V\|$	$\|E\|$	$\|\mathcal{Q}\|$	total demand
nsf14-d0	14	21	33	354
ge17-d0	17	26	58	740
uk24-d0	24	49	91	1302
eur42-d0	42	81	479	3336

Table 2: Characteristic parameters of the instances.

instance	used	total number of installed								
	links	S	L	XL	W16	W32	OXCS	OXCM	OXCL	conv.
nsf14-d0	18	2	10	19	2	29	12	2	0	3
nsf14-d50	21	8	14	19	8	33	11	3	0	21
nsf14-d0-pr	21	22	17	0	26	13	14	0	0	14
nsf14-d50-pr	21	22	21	2	23	22	13	1	0	21
ge17-d0	24	0	31	14	8	37	13	6	0	4
ge17-d50	26	3	43	15	8	53	15	5	1	12
ge17-d0-pr	26	23	32	1	29	27	14	5	0	10
ge17-d50-pr	26	28	43	3	34	40	16	6	0	0
uk24-d0	46	9	42	51	11	91	10	17	1	9
uk24-d50	49	3	57	67	8	119	5	20	2	42
uk24-d0-pr	49	48	59	21	63	65	11	18	0	1
uk24-d50-pr	49	51	77	30	65	93	7	21	1	19
eur42-d0	76	11	197	91	13	286	21	25	12	11
eur42-d50	81	26	227	134	30	357	12	30	16	15
eur42-d0-pr	81	87	207	47	95	246	14	25	13	6
eur42-d50-pr	81	101	252	89	113	329	20	33	13	22

Table 3: Computed network configurations.

5.2 Results

The computational results are summarized in two overviews. Table 3 displays the configuration of the network by the total number of links used and the total numbers of installed devices. The reported number of converters is achieved by the sequential application of four heuristics: three variants of the described greedy heuristic and the iterated local search method. For the last one, a time limit of 10% of the time spent so far is used.

Table 4 presents the number of lightpaths that has been rerouted in the postprocessing step of the dimensioning and routing phase, as well as the cost for the resulting networks together with the computed lower bound, the optimality gap, and the computation time in seconds. All computations have been carried out on a Pentium 4 processor with 1700 MHz and 1 GB RAM.

A first conclusion that can be drawn from these results is that the decomposition of the problem into a dimensioning and routing subproblem and a wavelength assignment subproblem seems to be reasonable from a computational point of view. The small number of required converters indicates that the separation of the wavelength assignment has at most a marginal impact on the total cost of the solution, whereas it is clear that computation times are reduced significantly. Concerning the wavelength assignment subproblem, none of the greedy heuristics turned out to be superior to the others. The iterated local search method has proven to outperform them on all instances, and hence is recommended for this task.

Table 3 shows that in the case of greenfield planning and the lack of a diversification requirement for the lightpath routing, devices are not installed on every link. These links however

instance	# re-routings	total costs	lower bound	gap	CPU time (s)
nsf14-d0	0	8692.10	7675.73	11.69%	54.0
nsf14-d50	12	11569.70	8374.53	27.62%	69.0
nsf14-d0-pr	0	2387.50	2153.59	9.80%	708.0
nsf14-d50-pr	6	3153.05	2526.53	19.87%	843.9
ge17-d0	0	4911.28	4539.18	7.58%	75.8
ge17-d50	3	6074.28	4799.32	20.99%	109.1
ge17-d0-pr	0	3961.65	3733.21	5.77%	911.4
ge17-d50-pr	0	5048.65	4092.57	18.94%	866.4
uk24-d0	0	9671.78	7416.31	23.32%	250.6
uk24-d50	32	11280.73	7703.12	31.71%	784.2
uk24-d0-pr	10	7994.22	5939.33	25.70%	2433.0
uk24-d50-pr	0	9852.86	6197.48	37.10%	2032.9
eur42-d0	0	37800.46	31048.30	17.86%	1502.8
eur42-d50	51	47320.41	33292.70	29.64%	12154.8
eur42-d0-pr	35	29516.68	23124.60	21.66%	1968.9
eur42-d50-pr	75	39987.37	24005.30	39.97%	6793.8

Table 4: Rerouting, total cost, lower bound, optimality gap, and computation time for each solution.

are needed in the case a survivable routing is desired. Notice also that the presence of pre-installed devices clearly leads to different link configurations, while the node configurations are quite similar. Although at each link two fibers are available for free, not all of them are used in all cases. This happens whenever less than two systems are needed in the solution.

Furthermore, Table 4 shows that the routings provided by DISCNET are almost always integer for practical instances. In half of the cases none of the lightpath routings is fractional, whereas for the other instances only a limited number of them has to be rerouted within the postprocessing. The additional cost arose by this procedure are typically very small. Again the computational effort is certainly reduced by this relaxation.

The lower bound enables us to evaluate the computed solutions. In this context, it should be noted that the lower bound is in fact composed of two parts, for the cost of the installed link and node technology, respectively. At present the computation of the lower bound on the OXC cost is less sophisticated than the one for the link technology. Since the node cost make up a considerable portion of the total cost in larger networks, this explains the in general larger gaps for these instances. Notice further that the lower bound does not imply that network configurations with cost equal or close to this bound exist, but guarantees that no cheaper solution can exist.

Both the gap and the total computation time show that the optimization problem is harder for the instances with 50% diversification. The costs of the solutions with survivability, however, show only limited increase in comparison to those without diversification. On the one hand, this indicates that there is space for improving the lower bound. On the other hand, by the relatively small gaps for the non-survivable instances, it is suggested that the computed solutions are rather good.

6 Concluding remarks

In this paper, a solution method is presented for transparent optical network design problems including three parts: the dimensioning of all network capacities, the routing of the lightpaths, and the wavelength assignment. The flexible representation of the network components within the model allows to incorporate many aspects which are important from a practical point of view. As a solution methodology, a decomposition approach has been favored that has two major advantages: the wavelength assignment is separated from the dimensioning and routing, and in addition to a solution, a lower bound on the total cost is provided. This approach has been shown to be appropriate for realistic problem instances.

A direction for further research could be the integration of additional aspects such as lightpaths operated on different bitrates. The application of other survivability concepts should also be considered. Moreover, advanced methods for integer multicommodity flows are desirable to improve the lightpath routing. Both from a theoretical and practical point of view, exact methods for solving the wavelength assignment subproblem are of interest. In particular, lower bounds on the number of unavoidable wavelength converters given a lightpath routing should be investigated.

References

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

[2] R. Dutta and G.N. Rouskas. A survey of virtual topology design algorithms for wavelength routed optical networks. *Optical Networks Magazine*, 1(1):73–89, January 2000.

[3] J. L. Kennington and Eli V. Olinick. Wavelength routing and assignment in a survivable WDM mesh network. *Operations Research*, 51(1), 2003. to appear.

[4] G. Li and R. Simha. On the wavelength assignment problem in multifiber WDM star and ring networks. *IEEE ACM Transactions on Networking*, 9(1):60–68, 2001.

[5] B. Mukherjee. *Optical Communication Networks*. McGraw-Hill, New York, 1997.

[6] R. Ramaswami and K.N. Sivarajan. *Optical Networks: A Practical Perspective*. Morgan Kaufmann Publishers, Inc., 1998. ISBN 1-55860-445-6.

[7] T.E. Stern and K. Bala. *Multiwavelength Optical Networks: A Layered Approach*. Addison Wesley Longman, Inc., 1999. ISBN 0-201-30967-X.

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

[9] H. Zang, J.P. Jue, and B. Mukherjee. A review of routing and wavelength assignment approaches for wavelength-routed optical wdm networks. *Optical Networks Magazine*, 1(1):47–60, January 2000.