Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ARIE M.C.A. KOSTER
ADRIAN ZYMOLKA

Minimum Converter Wavelength Assignment in All-Optical Networks

ZIB-Report 03–45 (December 2003)

Minimum Converter Wavelength Assignment in All-Optical Networks

Arie M. C. A. Koster* Adrian Zymolka*

Abstract

Finding conflict-free wavelength assignments with a minimum number of required conversions for a routing of the lightpaths is one of the important tasks within the design of all-optical networks. We consider this problem in multi-fiber networks with different types of WDM systems. We give a detailed description of the problem and derive its theoretical complexity. For practical application, we propose several sequential algorithms to compute appropriate wavelength assignments. We also perform computational experiments to evaluate their performance. For the iterative algorithms, we identify characteristic patterns of progression. Two of these algorithms qualify for application in practice.

1 Introduction

At the core of modern telecommunication networks, optics is used as transmission technology. By Wavelength Division Multiplexing (WDM), multiple signals can be transmitted through the same fiber, which enlarges its capacity significantly. With the deployment of Optical Cross Connects (OXCs), these signals can be switched at the nodes of the network without opto-electronic-optic (o-e-o) conversion. Such an ongoing optical signal is called a lightpath. In principle, the wavelength on which a lightpath is operated can be exchanged at every node along the path. Such an exchange however can only be performed if a wavelength converter is installed. The cost of such devices as well as the transmission delay caused by the conversion, are the reason for operators to minimize the number of conversions.

For the design and configuration of optical networks, three tasks have to be carried out:

- the dimensioning of the physical network by the installation of fibers and WDM systems as well as OXCs to provide sufficient transmission and switching capacities,
- the routing of lightpaths to satisfy the given connection demands, and

*Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), Takustraße 7, D–14195 Berlin, Germany. Email: {koster,zymolka}@zib.de

Keywords: minimum converter wavelength assignment, optical networks, generalized coloring, complexity, heuristics

Mathematics Subject Classification (2000): 90B18, 68M10, 90-08

- the assignment of wavelength to the lightpaths such that the number of required wavelength conversions is minimized.

In [14], this integrated network planning task is discussed, and a solution method is presented. The approach consists of the decomposition of the comprehensive optimization problem in the dimensioning and routing problem on the one hand and a subsequent wavelength assignment problem with converter placement on the other hand. Such a decomposition is possible since the availability of converters ensures the existence of a conflict-free wavelength assignment for any given dimensioning and routing solution. Such a decoupling of dimensioning, routing, and wavelength assignment is beneficial for two reasons: first of all, the complex design optimization for optical networks becomes computationally tractable (as shown in [14]). Second, as this study shows, wavelengths can be assignmed with very few conversions (accounting for a marginal contribution to the total cost) or even without conversion at all. Hence lower-cost solutions by integrated solution approaches will typically base on an improvement of the dimensioning and routing, not specifically of the wavelength assignment.

The problem to assign wavelengths to lightpaths with placement of converters has already been considered in the literature before. Here, we have to distinguish between devices that convert the wavelength of a single lightpath and those that perform conversion for all lightpaths crossing a node (by electronic switching devices instead of optical ones). The latter case is studied in for example [9, 12]. Most studies for the first case discuss the problem on specific topologies (e.g., [1, 13]), or other regular structures, cf. [5]. Moreover, the combined problem of routing and wavelength assignment has been extensively studied, see [3] for the case with wavelength converters and equal number of fibers at all links.

In this paper, we investigate the wavelength assignment problem including the placement of converters in a more general setting. We start with a problem description in Section 2 and discuss the computational complexity. For the calculation of feasible wavelength assignments, we present several algorithms in Section 3. A computational comparison of the various algorithms in Section 4 shows their performance on realistic instances. The results identify the methods which are especially suited for deployment in network design tools. Concluding remarks in Section 5 close the paper.

2 Minimum converter wavelength assignment

In this section, we describe the Minimum Converter Wavelength Assignment Problem (MCWAP) under consideration. We also discuss its computational complexity which shows that solving it to optimality is $\mathcal{NP}$-hard.

2.1 Problem description

Wavelength assignment is one of the tasks to be carried out during the design and configuration of an all-optical network. The overall objective is to design a cost-efficient optical network such that all traffic (lightpath requests) can be switched. Since this optimization problem is too complex to handle with state-of-the-art mathematical optimization tools, in [14] a decomposition approach is described which is implemented as software tool OND (Optical Network Design). In this approach, the wavelength assignment is done on top of a physical network configuration and a routing of the traffic demands. This in particular implies that:

- No restrictions on the network configuration are presumed, that is, at each link a number of fibers is operated with WDM systems of (in principle) arbitrary type, and at each node, one or multiple OXCs offer switching capability.
- The routing of the lightpaths is specified at the level of links and nodes, not at the level of fibers and OXCs.
- Multiple lightpaths can be routed between two nodes, using the same or different paths.
- The installed capacities are sufficient to accomplish the lightpath routing, i.e., the number of lightpaths using a link does not exceed the number of optical channels provided by the installed fibers and WDM systems on that link, and each node offers sufficient switching ports to handle all crossing lightpaths.

This general setting for the wavelength assignment problem is motivated by the optimization problem to be solved for the dimensioning and routing without distinguishing the wavelengths. In addition, the lightpaths are not dedicated to the fibers (and WDM systems), since only the total capacity of a link is of importance for their routing.

Although asymmetries in the physical equipment as well as in the lightpath routing can be simply incorporated, we restrict the discussions in this paper to symmetric transmission and switching capabilities, i.e., all available optical channels can be used in both directions, and the lightpaths are bidirectional as well. As a consequence, the physical network can be modeled as an undirected graph $\mathcal{N} = (N, L)$ with node set N and link set L. Undirected paths in $\mathcal{N}$ represent the lightpaths.

The configuration of the optical network is now completed by the determination of the wavelength of operation for each of the links of all lightpaths. Each WDM system provides a set of predefined wavelengths. Otherwise stated: each wavelength is available a limited number of times on each link. If, due to these capacity restrictions, the wavelength has to be changed along the path, a wavelength converter must be installed in the node of exchange. We consider only one type of wavelength converter that can translate the wavelength of a single lightpath from any wavelength to any other in the spectrum. We do not restrict the number of converters that can be installed in a node. However, each required wavelength converter causes a (node independent) cost. To minimize the overall network cost, the objective is to minimize the total cost for conversion, which is equivalent to minimizing the number of wavelength converters. So, the Minimum Converter Wavelength Assignment Problem (MCWAP) consists in the task to assign a wavelength to each link of each lightpath such that each wavelength is not used more often than available on each link, and the total number of wavelength conversions is minimized.

2.2 Complexity of the problem

It is well-known that wavelength assignment in optical networks is related to coloring the vertices of a graph. In case only a single fiber is installed at all links and the same WDM systems operated on all these fibers, the question whether a wavelength assignment without conversion exists is equivalent to the question whether a coloring of the so-called *path conflict graph* exists with the number of colors equal to the number of wavelengths. If such a coloring exists, it can easily be translated to a wavelength assignment without converters. However, if no such coloring exists, i.e., each coloring needs more colors than wavelengths available, it is unclear how such a coloring can be translated to a complete wavelength assignment. Moreover, in case the wavelengths are available more than once on the links, it is not necessary to assign different wavelengths to lightpaths that share the same link. Hence, the path conflict graph is of limited use for wavelength assignment in practical settings.

For the above reasons, we present in this section an alternative relation to coloring the edges of a graph. Given a graph $G = (V, E)$ without loops (while parallel edges are allowed), the *edge coloring problem* is to find a labeling of the edges with as few colors as possible such that all edges incident to a vertex are labeled with different colors. The *chromatic index* of a graph G, denoted by $\chi'(G)$, is the minimum number of colors in a feasible edge coloring. We first formally state the decision versions of both problems.

Chromatic Index [6]
instance: Graph $G = (V, E)$ and a positive integer K.
question: Does G have chromatic index K or less, i.e., can E be partitioned into $k \leq K$ pairwise disjoint sets $E_1, E_2, \ldots, E_k$ (representing the color classes) such that, for all $i = 1, \ldots, k$, no two edges in E_i share a common endpoint in G?

Minimum Converter Wavelength Assignment
instance: Let $\mathcal{N} = (N, L)$ be an undirected graph representing the optical network, and let Λ be the spectrum of all wavelengths. For each link $\ell \in L$, let κ_ℓ^λ be the number of times wavelength $\lambda \in \Lambda$ is available by the installed fibers and WDM systems. Let $\mathcal{P}$ be a collection of paths in $\mathcal{N}$ (not necessarily pairwise disjoint). Finally, let M be a non-negative integer.
question: Does there exist an assignment of the wavelengths $\lambda \in \Lambda$ to the links of the paths in $\mathcal{P}$ such that wavelength $\lambda \in \Lambda$ is assigned on link $\ell \in L$ at most κ_ℓ^λ times, and at most M converters are needed for all wavelength exchanges along the paths?

Holyer [8] proved that Chromatic Index is $\mathcal{NP}$-complete. Based on this result, we obtain for Minimum Converter Wavelength Assignment:

Theorem 2.1 Minimum Converter Wavelength Assignment *is $\mathcal{NP}$-complete, even if $\mathcal{N}$ is a star and $\kappa_\ell^\lambda = 1$ for all $\ell \in L$ and $\lambda \in \Lambda$.*

Proof. The theorem is proved by a reduction of Chromatic Index to the stated special case of Minimum Converter Wavelength Assignment.

Let a graph $G = (V, E)$ and a positive integer K define an instance of Chromatic Index. Then we construct a network $\mathcal{N} = (N, L)$ by setting $N = V \cup \{c\}$ and $L = \{vc | v \in V\}$.

So, $\mathcal{N}$ is a star with artificial center node c connected to all vertices $v \in V$. Furthermore, let the total spectrum $\Lambda = \{\lambda_1, \ldots, \lambda_K\}$ contain K wavelengths, and set the capacity for every link $\ell \in L$ and wavelength $\lambda \in \Lambda$ uniformely as $\kappa_\ell^\lambda = 1$. Finally, for every original

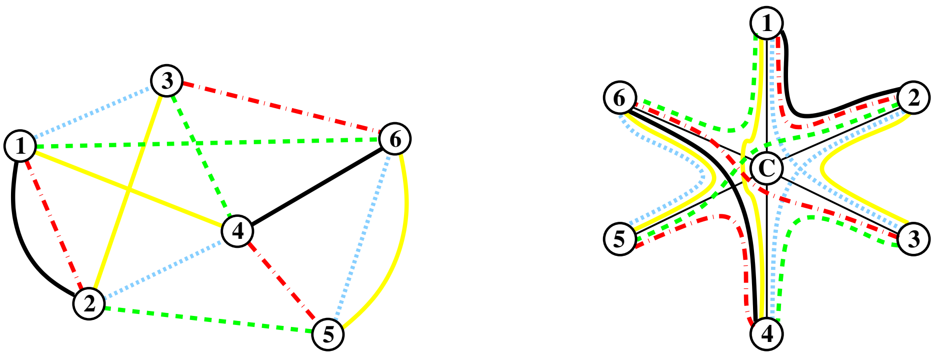

(a) Edge coloring instance

(b) Associated MCWAP instance

Figure 1: Reduction from Chromatic Index to Minimum Converter Wavelength Assignment

edge $vw \in E$, we add a lightpath $p_{vw} := [\{v, c\}, \{c, w\}]$ to $\mathcal{P}$ (cf. Figure 1).

Then G has chromatic index K or less if and only if there exists a wavelength assignment for the constructed Minimum Converter Wavelength Assignment instance with M=0 converters. Hence, solving Minimum Converter Wavelength Assignment even on this star network for the special case M=0 has at least the same complexity as edge coloring which is $\mathcal{NP}$-complete. ■

Note that the construction in the proof is reversible for the special case of star networks with a single fiber at all links. Let K be the number of available wavelengths. If an edge coloring of G is found with $L > K$ colors, the lightpaths corresponding to the $L - K$ least used color classes need a converter at c.

In contrast to the vertex coloring problem, the edge coloring problem can be generalized to incorporate the operation of multiple fibers and WDM systems on a link. If f fibers/WDM systems of the same type are installed on a link $vc \in L$, then at most f lightpaths sharing this link can be assigned the same wavelength. Such an assignment corresponds to a coloring of the edges in G such that at most f edges incident with v are labeled with the same color. This problem is known as the f-edge-coloring problem, see Hakimi and Kariv [7]. By other means, Li and Simha [11] showed that if all links are setup with k fibers/WDM systems and k is even, no wavelength conversion is necessary in the special case.

The $\mathcal{NP}$-hardness of the general MCWAP indicates that we cannot expect to find a polynomial time algorithm for computing optimal assignments. This thesis is confirmed by our computational experiments with an integer linear programming (ILP) formulation of the problem. The used model is straightforward and hence not presented explicitely here. Our computational experiments with this ILP have shown that realistic instances of MCWAP surpass the possibilities of state-of-the-art integer program solvers. Like with coloring problems, the symmetry between the wavelengths causes the LP relaxation to be fractional and typically attaining the value 0. Problem-specific branching strategies or cutting planes can not resolve this difficulty. From a practical perspective, it is therefore of importance to find algorithms that provide close-to-optimal assignments.

3 Sequential ordering algorithms

In this section, we discuss heuristic algorithms for MCWAP. We distinguish between constructive and iterative methods. Where constructive heuristics try to generate good assignments from scratch, the iterative approaches start with an assignment and try to reduce the converter number by clever transformations. Some of the heuristics presented in this section have been briefly discussed in [14], in context of the integrated optimization approach for optical network design. The common feature of all algorithms is that they base on the principle of sequential wavelength assignment.

3.1 Sequential wavelength assignment

The main idea for both the constructive and iterative heuristics is to assign wavelengths to the lightpaths in a sequential way, i.e., the lightpaths are processed one by one in a certain order. In each step, the locally best assignment is to assign the wavelengths to the lightpath at hand such that the number of required converters is minimized. Each time a lightpath is processed, the availability of the assigned wavelengths is reduced by one. Hence, even if we start with uniform wavelength capacities at the links, the remaining set of wavelengths differs from link to link during the process. A locally optimal assignment for a single lightpath can be found with the following procedure:

> Beginning with the first link, repeatedly select a wavelength that can be assigned as far as possible, until no links of the lightpath are left anymore.

It is easy to verify that this procedure minimizes the number of converters for the lightpath at hand: We start with a farthest reaching wavelength, thus any other assignment cannot place the first converter in a later node. So, all assignments need a converter before or in the same node. Since we continue again with a farthest reaching wavelength, the same argumentation applies (iteratively). As conclusion, any feasible assignment needs at least the same number of converters as the above procedure places.

To break ties in case multiple wavelengths reach equally far, we assume an ordering of the wavelengths. The first wavelength in this ordering that reaches farthest is selected.

The order in which the lightpaths are processed is of significant importance for the number of converters needed in the final wavelength assignment. As MCWAP is $\mathcal{NP}$-hard, a guaranteed optimal sequence cannot be expected to be determined in polynomial time. In fact, there need not to be such an ordering. Consider the (pathological) example in Figure 2 where two parts of a single all-optical network are displayed. In each part, four

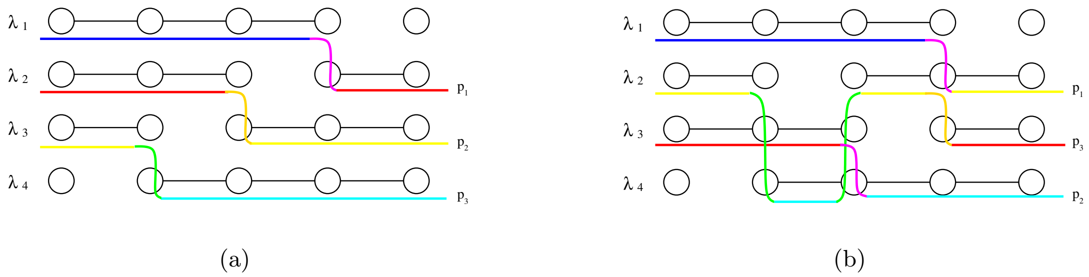

Figure 2: Example for an instance where a sequential heuristic finds an optimal wavelength assignment in (a), but not in (b) where only the spectrum has been slightly rearranged.

consecutive links are considered on which three lightpaths have to be established. Each row represents a certain wavelength which is available on a link if its endnodes are connected in this row. Note that both cases differ only in the order of the wavelengths. Moreover, thin lines illustrate the lightpaths constructed by using the farthest reaching wavelength method according to the wavelength ordering $(\lambda_1, \lambda_2, \lambda_3, \lambda_4)$. Changing a row means to place a converter at the appropriate node. While in the left part, only three converters are needed, the right solution requires five converters. Reordering the wavelengths to $(\lambda_1, \lambda_3, \lambda_2, \lambda_4)$ simply exchanges both sides. Each other order of the wavelengths results in a solution with at least the same number of converters or even more, while an optimal solution for the whole network clearly uses only six converters (applying different wavelength orders in the parts). This shows that there exist instances of the minimum converter wavelength assignment problem for which not all solutions can be generated by a sequential method. As a consequence, the space of all solutions that can be generated by a sequential procedure may be restricted to a subset of all feasible assignments. Especially, it is not sure that it contains an optimal solution, thus the procedure may detect a close-to-optimal solution at best.

On the other hand, however, sequential wavelength assignment has promising features as well, in particular if the optimal assignment only needs a small number of converters. If there exists a wavelength assignment without converters, then there also exists an ordering of the lightpaths such that the sequential wavelength assignment algorithm finds an optimal assignment. To see this, notice that the optimal assignment implies such an ordering: Arrange all lightpaths using the same wavelength consecutively. Then the sequential procedure finds the optimal assignment (without converters) if the wavelengths are assigned in the same order as they appear in the lightpath sequence.

In case the optimal solution needs a single converter, a similar argument applies. Consider the only lightpath that needs a converter. Order the wavelengths such that the wavelength used before conversion is the last but one, and the wavelength used after conversion is the last one. Now order the non-converted lightpaths using all but the last wavelength as before. Moreover, order the lightpaths operated on the last wavelength such that first those lightpaths are processed that contain the first link after conversion. By inserting the lightpath that needs conversion directly after these ones, the sequence provides a wavelength assignment that needs a single converter only.

This result can be further generalized to cases where only a few converters are needed in different parts of the network. So, at least if the number of required converters is expected to be small, sequential wavelength assignment seems to be an effective way to find close-to-optimal solutions.

3.2 Constructive methods

A very fast way to obtain a wavelength assignment is to define an ordering on the lightpaths in a greedy manner, i.e., we construct a solution by selecting the next lightpath to be processed according to a certain rule. Three ordering rules for the lightpaths have been implemented:

- longest path first,
- most inflexible path first, and
- most inflexible longest path first.

The *longest path first rule* selects among the not yet processed lightpaths one that contains the largest number of links. The idea behind this decision is that it becomes more difficult to assign wavelengths to long lightpaths without wavelength conversion if more wavelengths have already been assigned to other lightpaths. Ties are broken arbitrarily. Note that this ordering is static, i.e., it can be determined before the algorithm is executed.

The *most inflexible path first rule* selects among the not yet processed lightpaths one for which the number of continuing wavelengths is minimal. In case all remaining lightpaths still can get a wavelength assignment without conversion, lightpaths with scarce availability of common wavelengths at the complete path are selected first before conversion is necessary (due to other assignments). The idea is that those lightpaths are most likely to need converters if other assignments further reduce the set of available (ongoing) wavelengths. Lightpaths which need converters anyway are also processed first. This ordering is dynamic, since the next lightpath to be processed can first be determined after the assignment of previously processed lightpaths.

The *most inflexible longest path rule* combines the two previous rules. Among the most inflexible paths, a longest one is selected. In this way, the increased risk for the need of wavelength conversion on long paths with a scarcity of wavelengths is taken into account. As for the previous rule, this ordering is also dynamic.

Clearly, many other rules are possible. We have experienced that the more complicated the involved ideas are, the more time consuming the computation is, which directly restricts the attractivity of such rules. The possibility to reorder a sequence on the basis of the actual assignment is therefore of more interest for practice.

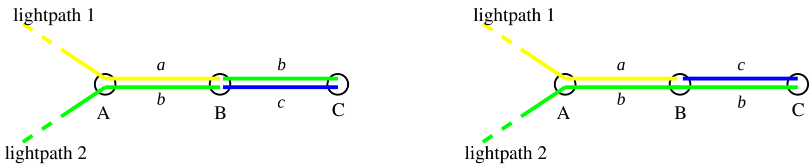

(a) 1 before 2 (b) 2 before 1

Figure 3: Basic concept for iterative of wavelength assignments

3.3 Iterative methods

Given a MCWAP solution, the number of needed converters can sometimes be reduced by an appropriate (partial) exchange of the wavelengths assigned to two lightpaths. Such an exchange can often be realized by interchanging the position of the lightpaths in the processing order. As example, consider the situation in Figure 3. There are two lightpaths with different beginnings, but both use the links AB, BC and end at node C. First suppose lightpath 1 is ordered before lightpath 2 and wavelength a is not available on link BC anymore, whereas wavelength c is not available on the links of lightpath 1 before node B. So, wavelength a is assigned to lightpath 1 on all links up to B, and wavelength b is used for link BC. Next, wavelengths are assigned to lightpath 2. Since wavelength b is not available on BC anymore, lightpath 2 gets wavelength b on the links up to B, but then needs to use wavelength c on link BC. This assignment is shown in Figure 3(a). However, it is clear that a reordering of the lightpaths allows for the reduction of the number of converters by one: Lightpath 2 gets wavelength b on all links, whereas lightpath 1 is assigned wavelength a up to node B and wavelength c on its last link BC. Figure 3(b) depicts this assignment (which results from that in Figure 3(a) only by exchanging the wavelengths on link BC).

This observation leads to the idea that exploiting the solution of a sequential method can offer information for a beneficial improvement of the processing sequence. We formalize this idea in an iterative method which base on a reordering of the lightpaths such that the number of converters needed for the new ordering is hopefully smaller than for the old ordering. The *iterative improvement procedure* works as follows. We start with an ordering of the lightpaths, either arbitrarily or initiated by one of the constructive heuristics. During the procedure, we keep track of the sequential order in which the lightpaths are processed. If an assignment without converters is found, we are done since this solution cannot be improved anymore. Otherwise, we choose one lightpath for which conversion is necessary, and put it at the beginning of the processing order. In this way, it most likely does not need a converter anymore (but may cause converters for other lightpaths). By iterating the sequential assignment algorithm and the reordering, better assignment are hopefully found. The best solution found during this process is stored.

Clearly, the described procedure represents a general method allowing for several variations. First of all, the selection of the lightpath to be pushed to the front allows for several strategies, for instance choosing the first lightpath(s) with converters, or the last, etc. Moreover, it is possible to vary the number of exchanges in each iteration. Instead of taking just one lightpath, two or more lightpaths with converters can be pushed at the beginning. For our computational experiments, we implemented the following four variants:

- push only the first lightpath that needs conversion to the beginning of the ordering,
- push only the last lightpath that needs conversion to the beginning of the ordering,
- push all lightpaths that need conversion to the beginning of the ordering, remaining the order of these lightpaths, and
- push all lightpaths that need conversion to the beginning of the ordering, reversing the order of these lightpaths.

It is clear that pushing (even single) lightpaths to the beginning of the ordering may yield a completely different result, a worse as well as a much better one. However, the first lightpaths in the ordering usually do not need conversions, so that at least the prepended lightpaths, which use converters in the former solution, have best chances to get rid of them. Hence, the iterative method explores many sequences, each started with a slight push in a (hopefully) promising direction. This effect may be accelerated by the last two variants which push many lightpaths at once to the beginning. The performed computational experiments serve to reveal this and other aspects of the proposed algorithms.

4 Computational experiments

In this section, we report on computational results obtained by the algorithms described in the previous section. We first describe how the instances are generated, and afterwards discuss the results.

4.1 Instance generation

Within the MultiTeraNet project of the German Ministry of Education and Research, the effort has been taken to propose three reference scenarios for optical transport networks, cf. [2]. These networks have 14, 17, and 28 nodes and describe respectively an US network based on the well-known NSF topology, a hypothetical German backbone network, and a pan-European network defined in the European COST 266 project. The traffic matrix for each of these networks has been generated by the model proposed in [4].

In total 24 instances are considered in this study, 8 for each network; all marked by `n` with the number of contained nodes. These instances are generated in the following way:

- For each network, we define four scenarios varying the fraction of the traffic that has to be protected against single node and link failures. Given a fraction $p \in [0, 1]$ and a demand of d lightpaths for a node pair, $\lceil pd \rceil$ lightpaths have to be protected. The used fractions are 0 (unprotected), 1/3, 2/3, and 1 (full protected). To indicate the applied level of protection, the instance name contains a `d` together with the number of thirds in the associated fraction, i.e., `d0` means unprotected demands, `d3` full protection, and `d1`, `d2` correspond to p being 1/3 and 2/3, respectively.

instance	$\|N\|$	$\|L\|$	p	P_1	P_2	P
n14-w1-d0/n14-w2-d0	14	21	0	2710	0	2710
n14-w1-d1/n14-w2-d1	14	21	1/3	2710	0	2710
n14-w1-d2/n14-w2-d2	14	21	2/3	2710	263	2973
n14-w1-d3/n14-w2-d3	14	21	1	2710	1668	4378
n17-w1-d0/n17-w2-d0	17	26	0	1021	0	1021
n17-w1-d1/n17-w2-d1	17	26	1/3	1021	1	1022
n17-w1-d2/n17-w2-d2	17	26	2/3	1021	319	1340
n17-w1-d3/n17-w2-d3	17	26	1	1021	873	1894
n28-w1-d0/n28-w2-d0	28	41	0	1008	0	1008
n28-w1-d1/n28-w2-d1	28	41	1/3	1008	140	1148
n28-w1-d2/n28-w2-d2	28	41	2/3	1008	499	1507
n28-w1-d3/n28-w2-d3	28	41	1	1008	885	1893

Table 1: Test instances for MCWAP, sorted by number of nodes and protection level.

- For each scenario with survivability requirements, we deploy the protection mechanism *Demand-wise Shared Protection* (DSP) proposed in [10]. DSP reduces the number of backup lightpaths needed for the required protection level by exploiting the connectivity of the physical topology to spread the lightpaths in the network.
- For each scenario, we consider two optical network design instances with different hardware capabilities. In the first instance, denoted by the annex `w1`, only a single type of WDM systems with 40 wavelengths, can be installed. In the second instance (`w2`), two different types of WDM systems, with 20 and 40 wavelengths respectively, are available. The underlying cost structure is realistic as reconciled with our industrial partners.
- With the optimization tool OND for optical network design [14], a dimensioning and routing is computed for all instances. The resulting hardware configuration and the lightpath routing serve as input for our wavelength assignment algorithms.

Table 1 summarizes some characteristics of the different instances. The values P_1, P_2 denote respectively the total traffic (in lightpaths) and the number of backup lightpaths needed by DSP to guarantee the protection level p. Hence, the total number of lightpaths is given by $P = P_1 + P_2$. Note that for instances `n14-w1-d1/n14-w2-d1` no backup lightpaths are necessary although it is guaranteed that 1/3 of the lightpaths survives a single node or link failure. By an appropriate diversification of the lightpath routing, this guarantee can be given. Note that the diversification (implying additional conditions) makes the routing for these instances different from the routings in the unprotected instances `n14-w1-d0/n14-w2-d0`.

4.2 Results

Both the constructive algorithms and the iterative methods have been implemented in C++ and run on the described instances. Since the iterative methods only stop if a solution without converters is found, a time limit of 600 CPU seconds (on a 2.53 GHz Intel Pentium 4 processor) has been included for a fair comparison between the iterative variants. In

instance	Constructive			Iterative			
	LPF	MIPF	MILPF	FPR	LPR	APR	APRR
n14-w1-d0	37	17	23	0	0	0	0
n14-w1-d1	105	152	109	71	103	0	0
n14-w1-d2	261	315	326	312	255	8	0
n14-w1-d3	352	395	562	543	485	33	0
n14-w2-d0	40	16	26	0	0	0	0
n14-w2-d1	109	135	116	64	102	0	0
n14-w2-d2	268	359	323	259	272	2	0
n14-w2-d3	296	343	529	508	425	0	0
n17-w1-d0	27	39	15	0	0	0	0
n17-w1-d1	41	55	51	0	0	0	0
n17-w1-d2	59	58	129	0	49	0	0
n17-w1-d3	132	234	223	142	147	12	9
n17-w2-d0	9	20	18	0	0	0	0
n17-w2-d1	54	53	43	0	0	0	0
n17-w2-d2	103	147	171	35	97	3	1
n17-w2-d3	124	279	179	102	160	13	6
n28-w1-d0	50	40	51	40	48	13	11
n28-w1-d1	12	11	9	0	0	0	0
n28-w1-d2	67	91	71	38	75	3	4
n28-w1-d3	84	213	135	52	112	0	0
n28-w2-d0	40	48	29	0	18	0	0
n28-w2-d1	47	66	60	32	76	21	19
n28-w2-d2	84	148	96	103	143	23	33
n28-w2-d3	194	229	246	184	213	50	59

Table 2: Results for the algorithms (LPF = longest path first, MIPF = most inflexible path first, MILPF = most inflexible longest path first, FPR = first path reordering, LPR = last path reordering, APR = all paths reordering, APRR = all paths reordering reversed).

Table 2 the results for all seven sequential algorithms are presented, whereas in Table 3 the number of iterations performed by the iterative methods are listed. Table 2 shows that the quality of the constructive heuristics is not comparable with that of the iterative methods. Among the constructive heuristics, no clear winner can be identified. Each of these three variants outclasses the other ones on certain instances.

More interesting is the situation among the iterative algorithms. Where first and last path reordering (FPR/LPR) only occasionaly find a solution without converters within the time limit, both other variants find zero value solutions on a more regular basis. As can be expected, moving a single path for reordering in each iteration of FPR and LPR provides a smaller potential progress on average. A significant speed-up can be obtained by reordering all paths with converters, either in the same (APR) or in reversed order (APRR). The APRR heuristic seems to be the most effective. Compared to APR, this might result from its more intensive permutational effect which brings additional motion into the generated lightpath orderings, thereby inspecting the search space in a larger scale. For 16 out of the 24 instances, it produces a wavelength assignment without conversions and finds this solution substantially faster than APR on several instances. For three cases, however, APR

instance	No. of Iterations			
	FPR	LPR	APR	APRR
n14-w1-d0	469	473	24	16
n14-w1-d1	1351	1354	455	110
n14-w1-d2	946	957	953	485
n14-w1-d3	637	646	637	277
n14-w2-d0	450	461	20	34
n14-w2-d1	1337	1344	142	126
n14-w2-d2	923	934	938	432
n14-w2-d3	645	650	412	389
n17-w1-d0	181	206	17	31
n17-w1-d1	517	1006	35	41
n17-w1-d2	1013	2015	157	70
n17-w1-d3	1368	1396	1365	1471
n17-w2-d0	218	300	27	29
n17-w2-d1	699	2264	139	205
n17-w2-d2	2070	2041	2167	2166
n17-w2-d3	1399	1430	1495	1517
n28-w1-d0	2714	2775	2811	2765
n28-w1-d1	293	392	27	43
n28-w1-d2	1317	1312	1349	1336
n28-w1-d3	1062	1060	993	574
n28-w2-d0	2215	2896	141	1220
n28-w2-d1	1842	1816	1837	1827
n28-w2-d2	1293	1291	1307	1306
n28-w2-d3	1001	1015	1018	1019

Table 3: Number of iterations for the iterative methods (FPR = first path reordering, LPR = last path reordering, APR = all paths reordering, APRR = all paths reordering reversed).

finds a better solution than APRR. The number of iterations (cf. Table 3) by the different iterative heuristics are comparable in case no solution without conversion is found.

Concerning the WDM system equipment, no difference in the optimal number of converters can be observed for the 14 node network. For the 28 node network, however, the instances with two types of WDM system seem to be harder than those with only one type.

A more detailed analysis of the performance of the iterative methods reveals that some typical progression patterns for the methods can be observed. These patterns can be put down to a couple of main characteristic behaviors as displayed in Figure 4. Each of the figures depicts the aquired number of converters in each iteration for all four methods. As first pattern, Figure 4(a) illustrates the frequent case in which all iterative algorithms find an optimal solution with none converters, but while APR and APRR terminate very fast, FPR and LPR take a much longer time. Figure 4(b) shows also a common behavior where APRR finds a zero value solution rapidly, APR approaches zero, but FPR and LPR seem to get stuck at a high number of converters.

Figure 4(c) demonstrates another typical pattern where APRR finds a best assignment first and APR after it. Here, FPR and LPR approach zero slowly, but steadily. Finally,

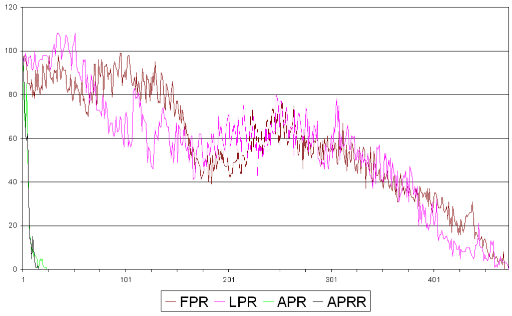

(a) n14-w1-d0

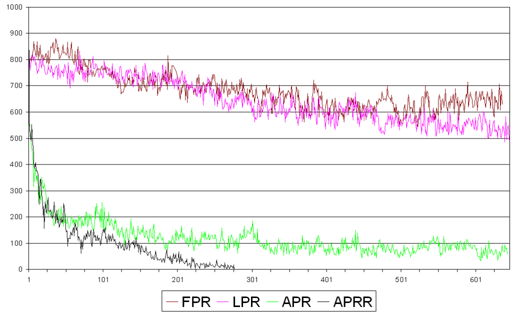

(b) n14-w1-d3

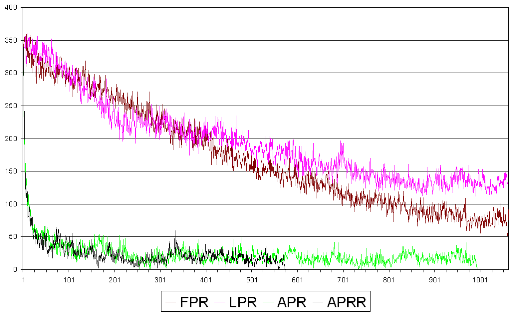

(c) n28-w1-d3

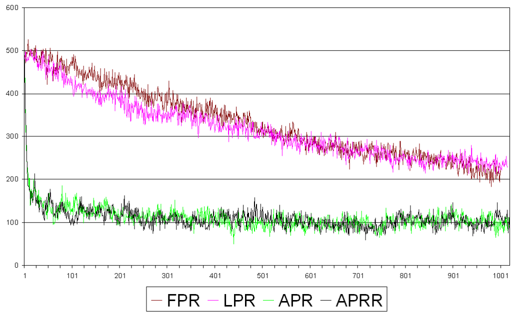

(d) n28-w2-d3

Figure 4: Number of converters needed in each iteration of the iterative methods for selected instances.

Figure 4(d) presents a case where both APR and APRR do not approach zero. In fact, the best solution found has been identified somewhere during the progress and afterwards only worse solutions with more converters are visited. FPR and LPR are still in the phase of decreasing the number of converters further. In order to see whether FPR and LPR would reach or even go below APR and APRR at some time, we ran these algorithms once again, now with one hour of computation time. The result is depicted in Figure 5 which unveils a remarkable effect: The number of converters by the LPR heuristic fluctuates heavily over time. Where FPR continues to approach the APR and APRR heuristic, the number of converters by LPR increases significantly at some point and then starts to yo-yo.

5 Concluding remarks

In this paper, we studied the Minimum Converter Wavelength Assignment Problem (MCWAP) in multi-fiber optical networks with different types of WDM systems. For the integrated planning of optical networks, the three tasks dimensioning, routing, and wavelength assignment have to be carried out. By the integration of wavelength converters, this complex

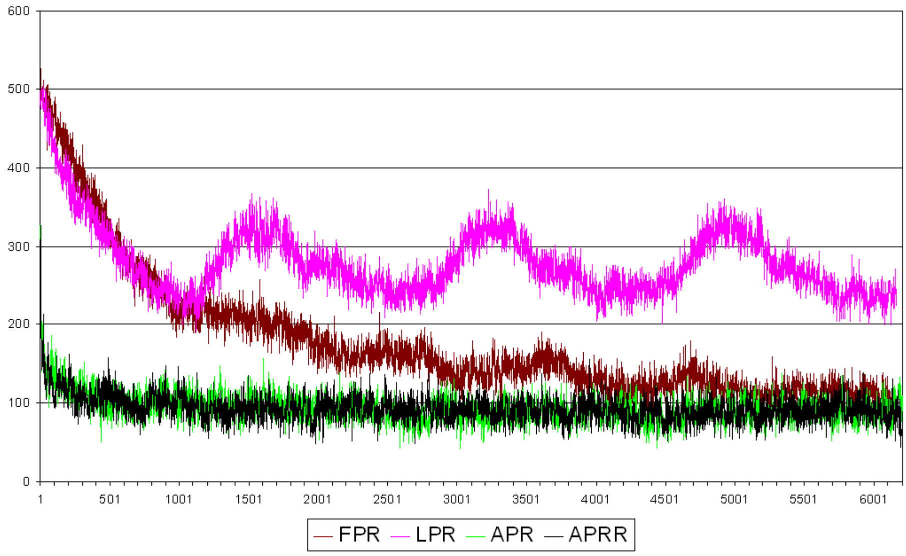

Figure 5: Number of converters needed in each iteration of the iterative methods for n28-d3-w2 with one hour of computation.

optimization problem can be decomposed in a dimensioning and routing subproblem and a wavelength assignment subproblem. Corresponding to the goal of cost-efficient network design, the objective of the wavelength assignment problem is to minimize the number of converters needed, given an appropriate dimensioning of the physical topology and a routing of the lightpaths.

After the specification of MCWAP, we discussed its complexity. A reduction from edge coloring exposed that MCWAP is $\mathcal{NP}$-hard even for simple star networks. For practical application, it is therefore useful to apply heuristic approaches. We described a series of both constructive and iterative methods and compared them by computational experiments. The results have shown that the all paths reordering algorithms, in particular the reversing variant, outperform all other ones. This way, feasible wavelength assignments without converters have been found for most instances. The fact that such assignments often exist and can be determined with the introduced algorithms confirms that the decomposition approach, e.g., decoupling dimensioning and routing and wavelength assignment, is valuable for a cost-efficient design of optical networks and suits for practical application.

As directions for further research, performance increasing improvements and alternative algorithms can be valuable. Furthermore, it is worthwile to derive lower bounds on the required number of converters. Such results are useful as guarantee for the quality of assignments computed by the sequential algorithms. Finally, an adaption of the algorithms to an environment with dynamic traffic requirements could be of interest .

References

[1] V. Auletta, I. Caragiannis, L. Gargano, C. Kaklamanis, and P. Persiano. Sparse and limited wavelength conversion in all-optical tree networks. *Theoretical Computer Science*, 266(1–2):887–934, 2001.

[2] A. Betker, C. Gerlach, R. Hülsermann, M. Jäger, M. Barry, S. Bodamer, J. Späth, C. Gauger, and M. Köhn. Reference transport network scenarios. Technical report, BMBF-Projekt MultiTeraNet, 2003.

[3] D. Coudert and H. Rivano. Lightpath assignment for multifibers WDM optical networks with wavelength translators. In *IEEE Globecom*, Taiwan, November 2002. OPNT-01-5.

[4] A. Dwivedi and R. E. Wagner. Traffic model for USA long-distance optical network. In *OFC 2000*, pages 156–158, 2000.

[5] A. Ferreira, S. Pérennes, A. W. Richa, H. Rivano, and N. Stier Moses. Models, Complexity and Algorithms for the Design of Multifiber WDM Networks. *Telecommunication Systems*, 2003. To be published.

[6] M. R. Garey and D. S. Johnson. *Computers and intractability: a guide to the Theory of $\mathcal{NP}$-Completeness*. Freeman and Company, N.Y., 1979.

[7] S. L. Hakimi and O. Kariv. On a generalization of edge-coloring in graphs. *Journal of Graph Theory*, 10:139–154, 1986.

[8] I. Holyer. The $\mathcal{NP}$-completeness of edge-coloring. *SIAM Journal on Computing*, 10:718–720, 1981.

[9] J. Kleinberg and A. Kumar. Wavelength conversion in optical networks. In *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 566–575, Baltimore, Maryland, 1999.

[10] A. M. C. A. Koster, A. Zymolka, M. Jäger, R. Hülsermann, and C. Gerlach. Demand-wise shared protection for meshed optical networks. In *Proc. of DRCN 2003*, page forthcoming, Banff, Canada, 2003.

[11] 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.

[12] G. Wilfong and P. Winkler. Ring routing and wavelength translation. In *Proc. 9th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 333–341, 1998.

[13] P. Winkler and L. Zhang. Wavelength assignment and generalized interval graph coloring. In *Proc. 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 830–831, 2003.

[14] A. Zymolka, A. M. C. A. Koster, and R. Wessäly. Transparent optical network design with sparse wavelength conversion. In *Proceedings of ONDM 2003*, pages 61–80, Budapest, Hungary, 2003. The 7th IFIP Working Conference on Optical Network Design & Modelling.