A post-optimization method to route scheduled lightpath demands with multiplicity

Lucile Belgacem

FuturMaster
696, rue Yves Kermen, 92100 Boulogne-Billancourt, France,
lucile.belgacem@futurmaster.com

Irène Charon

École nationale supérieure des télécommunications
46, rue Barrault, 75634 Paris Cedex 13, France
charon@enst.fr

Olivier Hudry

École nationale supérieure des télécommunications
46, rue Barrault, 75634 Paris Cedex 13, France
hudry@enst.fr

Abstract

We consider a NP-hard problem related to the routing and wavelength assignment (RWA) problem in optical networks, dealing with scheduled lightpath demands (SLDs) with multiplicity. A SLD is a connection demand between two nodes of the network, during a certain time. Each SLD requires a given number of wavelengths (its multiplicity). Given a set of SLDs, we want to assign a lightpath (i.e. a routing path and the necessary wavelengths) to each SLD, so that the total number of required wavelengths is minimized. To solve the problem, we design a post-optimization method allowing to improve the solutions provided by a heuristic.

Experimental results show that this post-optimization method is quite efficient to reduce the number of necessary wavelengths.

Keywords: WDM Optical Networks, Routing and Wavelength Assignment, Scheduled Lightpath Demands, Combinatorial Optimization, Post-Optimization

1 Introduction

We consider a problem related to the routing and wavelength assignment (RWA) problem in wavelength division multiplexing (WDM) optical networks. For a given network, represented by an undirected graph $\mathcal{G}$, the RWA problem consists in establishing a set of traffic demands S (or connection requests) in this network. Traffic demands may be of three types: static (permanent and known in advance), scheduled (requested for a given period of time) and dynamic (unexpected). A typical objective of RWA is to minimize the required number of wavelengths necessary to route all the demands.

This problem or variants of it have been extensively studied in the last decades (see, among others, [1], [2], [5], [7], [8], [9]). Many of these works consider static demands, without multiplicity. Here we deal with the case of a set S of scheduled lightpaths demands (SLDs) with multiplicity, which is relevant because of the predictable and periodic nature of the traffic load in real transport networks (more intense during working hours), but also much more difficult because of the time constraints which do not exist for static demands and because of the constraints involved by the multiplicities.

More precisely, a SLD s of S can be represented as $s = (x, y; \alpha, \beta; \nu)$, where x and y are some vertices of $\mathcal{G}$ (the source and destination nodes of the connection request), α and β denote the set-up and tear-down dates of the demand, and ν denotes the multiplicity of s. The routing of $s = (x, y; \alpha, \beta; \nu)$ consists in setting up a lightpath $(P; w_1, w_2, ..., w_\nu)$ between x and y, where P is a path (also called route) between x and y in $\mathcal{G}$ and w_i a wavelength for $1 \leq i \leq \nu$. In order to satisfy the SLD s, this lightpath (i.e. the path and the ν wavelengths) must be reserved during all the span of $[\alpha, \beta]$.

The wavelength continuity constraint is imposed: the same wavelengths must be used on all the links used by a lightpath and the path P must be the same for all the ν wavelengths. Moreover, at any given time, a wavelength can be used at most once on a given link; in other words, if two demands overlap in time, they can be assigned a same wavelength if and only if their routing paths are disjoint in edges.

We address the problem consisting in minimizing the number W of wavelengths required to establish all the SLDs. This problem is NP-hard, even if all the multiplicities are equal to 1 (see [3]). A solution to this problem is defined by specifying, for each SLD, the lightpath chosen for supporting the connection (i.e. a route and ν wavelengths), so that there is no conflict between any two lightpaths. Several approximate or exact methods have been proposed to deal with this NP-hard problem for static demands or for SLDs (see for instance [2], [4], [6], [8] and [9]).

The greedy method proposed by N. Skorin-Kapov in [8] gives very satisfying results in a very small amount of time and is, with this respect, among the most efficient heuristics. Its application will be used as a benchmark for measuring the performance of our post-optimization method.

2 The Post-optimization Method

The post-optimization method aims at improving the results provided by any heuristic permitting to solve the addressed problem. It consists in minimizing the overall values of the wavelengths of the established lightpaths in order to try to minimize the total number W of wavelengths. In the following, the wavelengths will be called 1, 2, ..., W.

Broadly speaking, the aim of the post-optimization method is to empty iteratively the set of SLDs of which the routing involves W in order to make W decrease successively. The principle of the method is the following: let $s = (x, y; \alpha, \beta; \nu)$ be any SLD, and let $(P; w_1, w_2, ..., w_\nu)$ be the current lightpath associated with s. Then we try to assign a new lightpath $(P'; w'_1, w'_2, ..., w'_\nu)$ to s, with $w'_\nu < w_\nu$. This is done by trying to assign to s the wavelengths $(1, 2, ..., \nu)$, then $(2, 3, ..., \nu + 1)$, $(3, 4, ..., \nu + 2)$, ..., until we find a possibility to reroute s or until we try $(w_\nu - \nu, w_\nu - \nu + 1, ..., w_\nu - 1)$ in vain.

More precisely, let us assume that we want to assign the wavelengths $(i, i+1, ..., i + \nu - 1)$ to s for some i with $i + \nu - 1 < w_\nu$. It is very likely that some of the SLDs s' currently using at least one of the wavelengths i, $i + 1$, ..., or $i + \nu - 1$ prevent us from routing s with these wavelengths. In other words, if we delete from $\mathcal{G}$ all the edges used to establish these SLDs s' which overlap s in time, we may find no path joining x and y. So, when we try the wavelengths $(i, i + 1, ..., i + \nu - 1)$, we consider a graph $\mathcal{H}(s)$, initially equal to $\mathcal{G}$, and we examine one after the other the demands s' which overlap s in time and which uses at least one of the wavelengths $i, i + 1, ..., i + \nu - 1$. For each such s', we remove from $\mathcal{H}(s)$ the edges of the path $P_{s'}$ supporting the connection associated with s' which are still in $\mathcal{H}(s)$. If there still exists a path in $\mathcal{H}(s)$ to set up s, we move up to the next conflicting demand of type s'; otherwise we cancel the lightpath associated with s', we put aside s' in a set E, and we put back the removed edges of $P_{s'}$ inside $\mathcal{H}(s)$ (of course, if some edges of $P_{s'}$ had been removed previously from $\mathcal{H}(s)$ because of former clashing SLDs, they remain removed). Thus, once all conflicting demands s' have been examined, it becomes possible to route s using the wavelengths $(i, i+1, ..., i+\nu-1)$ since all the conflicting lightpaths have been (at least temporarily) removed.

We must now deal with the demands of E. For each SLD s' belonging to E, we try to reroute s' without modifying the routing of any other SLD and only with the wavelengths $1, 2, ..., w_\nu - 1$ (where w_ν was the maximum wavelength used by s). In other words, we try to select ν' wavelengths among $1, 2, ..., w_\nu - 1$ (we do not require that the selected wavelenghts define an interval) such that it is possible to find a path in the graph $\mathcal{H}(s')$, defined in a similar way as for the graph $\mathcal{H}(s)$ above. If such a combinaison can be found for each demand s' of E, then we have finished with the demand s: we keep the tried wavelengths $(i, i+1, ..., i+\nu-1)$ for s, with a path disjoint with the ones of the other SLDs using at least one of the wavelengths $(i, i+1, ..., i+\nu-1)$, and we move up to another demand s to which we apply the same process. Otherwise we consider that the attempt to assign the wavelengths $i, i+1, ..., i+\nu-1$ to s has failed: we restore the original lightpaths of the SLDs belonging to E, and we try to assign the wavelengths $i+1, i+2, ..., i+\nu$ to s by the same process. If all the intervals $(i, i+1, ..., i+\nu-1)$ have been tried in vain, s remains with its original lightpath, and we move up to another demand s to which we apply the same process.

Thanks to this rearrangement, we may from time to time reroute all the SLDs using the wavelength W. In this case, we definitely remove this wavelength: we have saved one wavelength and W becomes $W-1$.

This method is referred to as the *post-optimization algorithm*. The overall heuristic consisting of the greedy algorithm followed by the application of the post-optimization method will be denoted $G+$ in the sequel.

3 Experiments

3.1 Framework

Due to the lack of space, we present here only some results obtained for two graphs, called $\mathcal{G}29$ and $\mathcal{G}57$. The first one, with 29 vertices and 44 edges, represents a hypothetical North-American backbone network; the second one, with 57 vertices and 85 edges, is extracted from the European optical transport network. We then randomly generate different sets of SLDs, with 500 or 1000

SLDs requiring up to 10 wavelengths each, leading to three types of instances, called $\mathcal{G}29$-500, $\mathcal{G}29$-1000 and $\mathcal{G}57$-500: $\mathcal{G}n$-m means that we consider the graph $\mathcal{G}n$ with m SLDs. For these sets of SLDs, the number of time-overlaps is significant but not too large: if they are too few, there are few clashes between the demands and therefore the addressed problem becomes too easy; on the contrary, if the time-overlaps are too numerous, the number of required wavelengths increases greatly and the problem becomes again less interesting.

In the following, we present the results obtained when applying the three heuristics G, RG and $G+$ to these three instances. Other experiments have been done, on different types of graphs and different sets of SLDs, leading to the same qualitative conclusions. The experiments have been performed on Solaris Sun stations (Sun Ultra 20M2 AMD bicore 3 Ghz).

3.2 Results

The results obtained for the three types of instances and the three heuristics are given in Table 1. For each case, we specify the average of the required numbers of wavelengths over 100 runs as well as the average CPU time in seconds; the last two lines specify the ratios $(W_G - W_{G+})/W_G$ and $(W_{RG} - W_{G+})/W_{RG}$, where W_G, W_{RG} and W_{G+} denote the average numbers of wavelengths obtained by applying 100 runs of G, RG and $G+$ respectively.

	$\mathcal{G}29$-500	$\mathcal{G}29$-1000	$\mathcal{G}57$-500
G	60.92 - 0.0091 sec	93.41 - 0.036 sec	80.42 - 0.015 sec
RG	51.32 - 20.55 sec	85.98 - 180.9 sec	71.72 - 62.30 sec
$G+$	46.68 - 19.25 sec	78.50 - 178.8 sec	65.65 - 61.59 sec
$(W_G - W_{G+})/W_G$	23.37 %	15.96 %	18.37 %
$(W_{RG} - W_{G+})/W_{RG}$	9.04 %	8.70 %	8.46 %

Table 1
Average numbers of required wavelengths and average CPU times

The results of Table 1 shows that, for these instances, $G+$ provides better results than G and that this remains true even if we give the same CPU time to both methods: $G+$ still provides better results than RG. This phenomenon is in fact quite general: in our experiments, with a same CPU time, $G+$ always provided better results than RG (and of course than G); in other words, the *gains* $(W_G - W_{G+})$ and $(W_{RG} - W_{G+})$ are never negative. Figures 1, 2 and 3

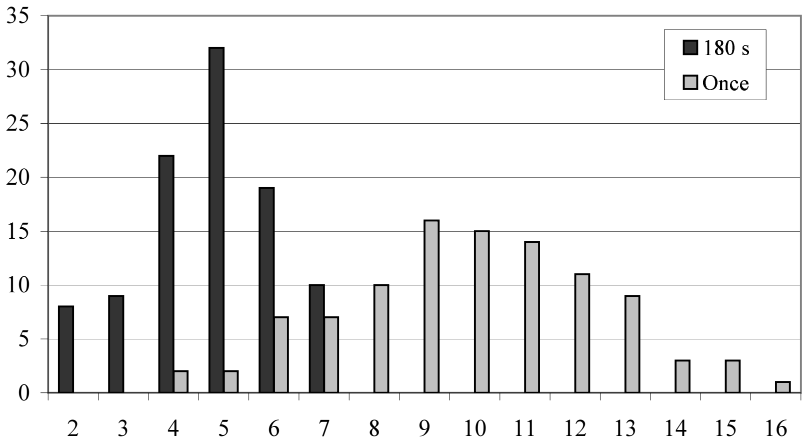

Fig. 1. Distributions of W for $\mathcal{G}$29-500

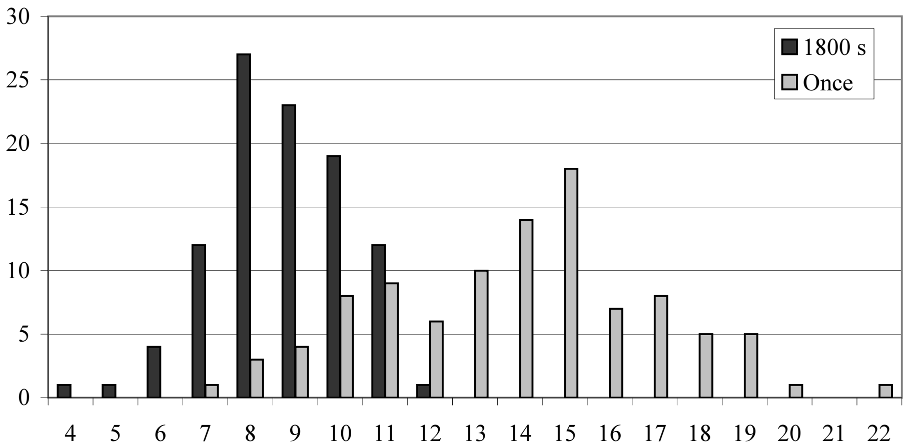

Fig. 2. Distributions of W for $\mathcal{G}$29-1000

display the distributions of these gains over 100 runs applied respectively to $\mathcal{G}$29-500, $\mathcal{G}$29-1000 and $\mathcal{G}$57-500. In each figure, the x-axis represents the value of $W_G - W_{G+}$ when G and $G+$ are applied once (on the right) or of $W_{RG} - W_{G+}$ when G and $G+$ are repeated (on the left; the CPU time devoted to each method is specified in the figures: 180 seconds for 500 SLDs, 1800 seconds for 1000 SLDs); the y-axis represents the number of times that each value of the gain has been observed.

According to these results (as said above, other instances have been studied, with other networks and other sets of SLDs, and the same type of results have been obtained every time), the post-optimization method appears as improving significantly the results given by the greedy heuristic, which were

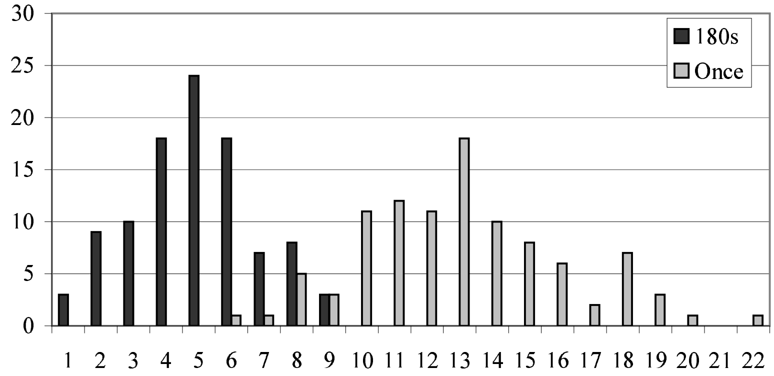

Fig. 3. Distributions of W for $\mathcal{G}57$-500

already good, while it is known in combinatorial optimization that reducing the gap between the computed solutions and the optimal ones becomes more and more difficult when going closer to the optimum.

For the instances considered here, the gain yielded by the application of $G+$ with respect to the sole application of the greedy heuristic G (more precisely the ratio $(W_G - W_{G+})/W_G$) is quite significant: between 16 % and 23 % for the results reported here. Even when considering the same CPU time, the gain of $G+$ with respect to RG (measured similarly by $(W_{RG} - W_{G+})/W_{RG}$) remains significant: between 8 % and 9 %.

Another important asset of $G+$ should also be mentioned: $G+$ succeeds in finding some values of W that neither G nor RG can reach during the 100 runs. When the number of SLDs increase, the gap between the values computed by $G+$ on the one hand and those computed by G and RG on the other hand becomes larger. Moreover for heavy loads of traffic demands (1000 or 3000 according to the considered network), the histograms for $G+$ become apart completely from the ones of G and RG: the worst solution provided by $G+$ remains better than the best solution found by G or RG.

On the other hand, $G+$ is significantly longer than G. In our experiments, the CPU time required to perform $G+$ can reach few minutes, whereas it is nearly instantaneous for G (less than one second). From a practical point of view, these computation times remain quite acceptable (especially considering the high complexity of the problem and the large sizes of the instances) since the addressed problem concerns connection requests that are known in advance. Indeed, in this case, a telecommunications operator can easily afford to spend the time required by the application of the post-optimization method in order to save some wavelengths, that will be available to establish further connection requests (unexpected demands for instance).

We may conclude that the post-optimization method improves the greedy algorithm significantly and in a reasonable CPU time. Even if G is repeated, it remains clearly better to use $G+$ than RG. Moreover this method can be applied to any other heuristics to deal with the RWA of SLDs, and even to other problems related to RWA in optical transport networks. It will be the topic of our next studies.

References

[1] Banerjee D., and B. Mukherjee, *A practical approach for Routing and Wavelength Assignment in Large Wavelength-Routed Optical Networks*, IEEE Journal on Selected Areas in Communications, 14 (5), 903–908, 1996.

[2] Belgacem L., and N. Puech, "Solving Large Size Instances of the RWA Problem Using Graph Partitioning", Proceedings of the 12th Conference on Optical Network Design and Modelling, Vilanova i la Geltru, Spain, 1–6, 2008.

[3] Chlamtac I., A. Ganz, and G. Karmi, *Lightpath communications: an approach to high-bandwidth optical WANs*, IEEE Trans. Commun. 40, 1171–1182, 1992.

[4] Gi Ahn H., T.-J. Lee, M. Y. Chung, and H. Choo, "RWA on Scheduled Lightpath Demands in WDM Optical Transport Networks with Time Disjoint Paths", Proc. International Conference on Information Networking, 342-351, 2005.

[5] Kuri J., N. Puech, M. Gagnaire, E. Dotaro, and R. Douville, *Routing and Wavelength Assignment of Scheduled Lightpaths Demands*, IEEE Journal on Selected Areas in Communications 21 (8), 1231–1240, 2003.

[6] Park S., J. S. Yang, and Y.-C. Bang, *On RWA Algorithms for Scheduled Lightpath Demands*, International Journal of Computer Science and Network Security 7 (3), 144-150, 2007.

[7] Ramaswami R., and K.N. Sivarajan, *Routing and Wavelength Assignment in All-Optical Networks*, IEEE/ACM Transactions on networking 3 (5), 489–500, 1995.

[8] Skorin-Kapov N., *Heuristic Algorithms for the Routing and Wavelength Assignment of Scheduled Lightpath Demands in Optical Networks*, IEEE Journal on Selected Areas in Communications 24 (8), 2-15, 2006.

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