Minimizing the number of wavelengths for the Routing and Wavelength Assignment Problem in optical network

L. Belgacem[1], I. Charon[2] and O. Hudry[2]
[1] Futur Master, Boulogne-Billancourt, France
[2] École nationale supérieure des télécommunications, Paris, France

***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). A SLD is a connection demand between two nodes of the network, during a certain time. Given a set of SLDs, we want to assign a lightpath (i.e. a routing path and a wavelength) to each SLD, so that the total number of required wavelengths is minimized. The constraints are the following: a same wavelength must be assigned all along the edges of the routing path of any SLD; at any time, a given wavelength on a given edge of the network cannot be used to satisfy more than one SLD. To solve this problem, we design a post-optimization method allowing to improve the solutions provided by a heuristic. The experimental results show that this post-optimization method is quite efficient to reduce the number of necessary wavelengths.**

***Index Terms*—Combinatorial Optimization, Heuristics, Post-Optimization, Routing and Wavelength Assignment Problem, Telecommunications, Wavelength Division Multiplexing Optical Networks.**

I. INTRODUCTION

We consider a problem related to the routing and wavelength assignment (RWA) problem in wavelength division multiplexing (WDM) optical networks (see e.g. [6], [7], [14], [20] or [29] for general references). For a given network topology, represented by an undirected graph G, the RWA problem consists in establishing a set of traffic demands S (or connection requests) in this network. Different versions of the RWA problem can be found in the literature, depending on the performance metrics and on the traffic assumptions (see for instance [28]). Traffic demands may be of three types: static (permanent and known in advance), scheduled (requested for a given period of time) and dynamic (unexpected). The typical objectives of RWA can be:

- to minimize the required number of wavelengths under given connection requests,
- to minimize the blocking probability, i.e. the number of rejected traffic demands, under given number of wavelengths and static or dynamic connection requests,
- to minimize the maximum number of wavelengths going through a single fibre,
- to minimize the network load as defined by the fraction of the number of wavelengths used on the overall set of fibre links in the network.

These problems have been extensively studied in the last decades (see, among others, [1], [2], [3], [5], [9], [10], [11], [12], [13], [14], [15], [16], [17], [19], [23], [28], [29]). Many of these works consider static demands. In our work, we deal with the case of a set S of scheduled lightpaths demands (SLDs), which is relevant because of the predictable and periodic nature of the traffic load in real transport networks (more intense during working hours, see [14]), but also much more difficult because of the time constraints which do not exist for static demands.

More precisely, an SLD s belonging to S can be represented by a quadruplet $s = (x, y, \alpha, \beta)$, where x and y are some vertices of G (source and destination nodes of the connection request), and where α and β denote the set-up and tear-down dates of the demand). The routing of $s = (x, y, \alpha, \beta)$ consists in setting up a lightpath (P, w) between x and y, where P is a path (also called route) between x and y in G and w a wavelength. In order to satisfy the SLD s, this lightpath must be reserved during all the span of $[\alpha,\beta]$.

The constraints related to the use of an optical network are the following:

- the same wavelength must be used on all the links used by a lightpath (wavelength continuity constraint);
- 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 the same wavelength if and only if their routing paths are disjoint in edges (wavelength clash constraint).

We address here the problem consisting in minimizing the number of wavelengths W required to establish all the SLDs. This problem is NP-hard (see [4]). A solution of this problem is defined by specifying, for each SLD, the lightpath chosen for supporting the connection (i.e. a route and a wavelength), so that there is no conflict between any two lightpaths. Several approximate or exact methods have been proposed in the literature to deal with this NP-hard problem for static demands or for SLDs: some are based on metaheuristics such as Tabu search or simulated annealing, others are based on greedy approaches more or less sophisticated (see [8], [9], [14], [15], [16], [18], [21], [22], [24], [25], [26], [27] and [28]).

The greedy method proposed by N. Skorin-Kapov in [23], which has been designed initially to deal with the case where SLDs may require several lightpaths simultaneously, gives very satisfying results in a very small amount of time and is, with this respect, among the most efficient heuristics. Its application to our problem (see below) will be used as a benchmark for measuring the performance of our approach. Indeed we propose in this paper a post-optimization method in order to improve the results given by other heuristics. The CPU time of the overall method will naturally increase, but it will remain acceptable to deal with SLDs: since the demands are known in advance, the allotted time to provide a solution is relatively large (unlike the case where demands are unexpected, and for which routings must be computed dynamically).

The greedy algorithm derived from [23] and the post-optimization method are presented in Section II. In Section III, we apply these methods on the American and European backbone networks, and on a large network generated randomly. For each case, we consider heavy traffic loads consisting of 500, 1000, and 3000 SLDs. Finally we analyse the obtained results and conclude in Section IV.

II. RESOLUTION METHODS

We describe in this section the different methods that we apply to solve RWA. We first present the greedy algorithm derived from [23]; we propose a slight modification of this algorithm so that it can be repeated. The post-optimization method is then described.

A. The Greedy Algorithm

The greedy algorithm derived from [23] consists in considering the wavelengths one by one, and in trying to route as many SLDs as possible with each wavelength. More precisely, let w be the current wavelength and $s = (x, y, \alpha, \beta)$ be the current SLD. We consider a graph $H(s)$ obtained from G by removing all the edges unavailable for the routing of s with w, i.e. edges that are contained in lightpaths corresponding to SLDs already routed with the wavelength w and which overlap s in time. According to this construction, any edge of $H(s)$ could be used to route the demand s with wavelength w without inducing any clash with previously established SLDs.

Thus, if there exists at least one path between x and y in $H(s)$, we attribute the shortest possible path P_S to the SLD s as well as the wavelength w; otherwise s is put aside and will be dealt with latter using another wavelength. Then we move up to the next not yet established SLD.

When all the SLDs have been examined, we move up to the wavelength $w + 1$ and try to route the remaining SLDs. The algorithm stops as soon as all the SLDs have been established: the current value of w specifies the sought value of W.

This algorithm will be referred to as *Gr*.

B. The Post-optimization Method

The post-optimization method presented in this paper aims at improving the results provided by *Gr*, though it can be applied to any heuristic designed 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 of wavelengths W.

The principle of the method is the following: for any $w \in \{2, ..., W\}$, we try to empty the set of SLDs routed with w, at least partially; this set will be called the *layer w* in the following. This is done by trying to assign a smaller wavelength $(1,2, ..., w - 1)$ to the demands of the layer w, which leads us to rearrange the wavelengths assigned to the SLDs of these lower layers. During this operation, the layers of some SLDs may change but all of them must remain in $[1, w - 1]$.

More precisely, let us assume that we want to move the demand $s = (x, y, \alpha, \beta)$ from its current layer w to a lower layer l ($l \in [1, w - 1]$). It is very likely that some of the demands belonging to the layer l prevent us from routing s with this wavelength. In other words, if we delete from G all the edges used to establish the demands of this layer which overlap s in time, we may find no path joining x and y.

So we consider a graph $H(s)$, initially equal to G, and we examine one after the other the demands s' of the layer l which overlap s in time. For each such s', we remove from $H(s)$ the edges of the path $P_{s'}$ supporting the connection associated with s' which are still in $H(s)$. If there still exists a path in $H(s)$ to set up s, we move up to the next demand s' of the layer l; otherwise s' is removed from the layer l and put aside in a set E, and we put back the removed edges of $P_{s'}$ inside $H(s)$ (of course, if some edges of $P_{s'}$ had been removed previously from $H(s)$ because of former clashing SLDs, they remain removed). Thus, once all demands of the layer l have been examined, it becomes possible to route s using the wavelength l since all the conflicting lightpaths have been (at least temporarily) removed.

We must now deal with the demands of E: we try to place each of these demands in one of the layers 1, ..., $w - 1$, without modifying the routing of any other SLD. If a layer can be found for each demand of E, then we have finished with the demand s: s remains in the layer l (and the demands of E remain in their new layers), with a lightpath compatible with the ones of the other SLDs of this layer, and we move up to the following demand of the layer w. Otherwise we consider that the attempt to move s to the layer l has failed: layer l is restored as before the attempt to insert s inside, and we try to move s to the next layer $l + 1$. If all the layers from 1 to $w - 1$ have been examined in vain, s remains inside its current layer w, and we move up to the following demand of the layer w.

As all the demands of the current layer w are handled, the number of remaining demands on this layer may decrease, and the layer w may become totally empty. In this case, we shift the layers $w + 1$, $w + 2$, ..., W to the layers w, $w + 1$, ..., $W - 1$, and we have succeeded in saving one wavelength definitively: W becomes $W - 1$.

This method is referred to as the *post-optimization algorithm*. Let us notice that even if a rearrangement of the layers does not permit to decrease the number W of required wavelengths, it may happen that further applications of the algorithm succeed to do so, because the SLDs are not dispatched in the layers in the same manner from one application to another. In the experiments presented below, we chose to repeat the post-optimization algorithm until four consecutive runs do not decrease *W*. This choice is based on an experimental observation and arises from a compromise between CPU time and the quality of the computed solutions. The overall heuristic consisting of the greedy algorithm followed by the application of the post-optimization method will be denoted *Gr+* in the sequel.

C. Repetition of Gr

The introduction of the post-optimization method yields a significant increase in computation time. To evaluate the post-optimization method, we will compare the results provided by the greedy heuristic with or without this post-optimization method. In order to avoid any bias, it is desirable that both methods are given the same amount of CPU time to provide a solution.

Thus we propose a slight modification of *Gr* in order to make it stochastic. Then we will be able to repeat *Gr* profitably as many times as necessary to attain the same CPU time as the one required by *Gr+*.

The principle is straightforward: we propose to consider, for each run of *Gr*, a random order for the examination of the SLDs. In [23], the demands are ordered with respect to the decreasing numbers of connection requests, which is irrelevant in our context since all the SLDs are assumed to require the establishment of only one lightpath.

According to our experiments, the best way to take benefit from the allotted CPU time seems to generate a random order of the SLDs for each run of *Gr*. Of course, the solution returned by this repetition of *Gr* will be the best one computed over the different runs of *Gr* during this repetition. This heuristic will be denoted *RGr*.

III. EXPERIMENTS

A. Framework

We present the results obtained for three graphs:

- G57 (57 vertices and 85 edges), extracted from the European optical transport network;

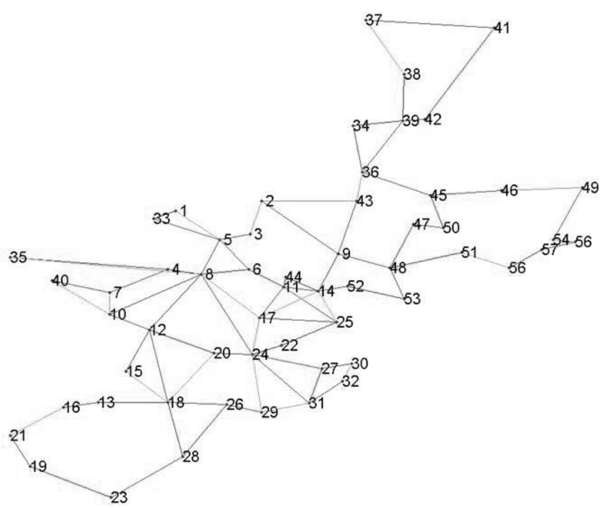

Figure 1. G57

- G29 (29 vertices and 44 edges), representing a hypothetical North-American backbone network;
- G200 (200 vertices and 239 edges), simulating a large optical network.

The first two graphs (see Figures 1 and 2) are often used to illustrate RWA problems. We have added the graph G200 in order to observe the behaviour of the methods when applied to larger networks. This graph has relatively few edges so that the number of paths that may support a given demand is limited (otherwise the number of required wavelengths is very small).

The sets of SLDs are generated randomly so that 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. Three sets of demands have been generated for each considered network with respectively 500, 1000 and 3000 SLDs. Thus we obtain nine instances. The name of each one is obtained by the concatenation of the name of the network with the number of SLDs to route. For instance G57-500 denotes the instance for which the network is G57, with 500 SLDs.

In the following, we present the results obtained when applying the three heuristics described above (*Gr*, *RGr* and *G+*) to these nine instances. The experiments have been performed on Solaris Sun stations (Sun Ultra 20M2 AMD bicore 3 Ghz). In order to evaluate the three heuristics on these instances, we carry out 100 runs of each method for each instance.

B. Results

The results obtained for the nine instances and the three heuristics are given in Tables 1 to 3 below. For each case, we specify the average of the required numbers of wavelengths over the 100 runs as well as the average CPU time in seconds; the last two lines of each table specify the ratios $\rho(Gr, Gr+) = (W_{Gr} - W_{Gr+})/W_{Gr}$ and $\rho(RGr, Gr+) = (W_{RGr} - W_{Gr+})/W_{RGr}$, where W_{Gr}, W_{RGr} and W_{Gr+} denote the average numbers of wavelengths required to set up all the connections when applying 100 runs of *Gr*, *RGr* and *Gr+* respectively.

Let us recall that the method *RGr* consists in repeating the greedy algorithm *Gr* as many times as required to attain the same computation time as *Gr+*. For example, for the instance G57-500, *Gr* runs in 0.02761 seconds whereas Gr+ takes 2.68 seconds. Therefore *RGr* is obtained by repeating *Gr* 97 times, which indeed corresponds to an overall computation time of 2.68 seconds. As said above, the result provided by *RGr* is of course the smallest number of wavelengths obtained during the repetition.

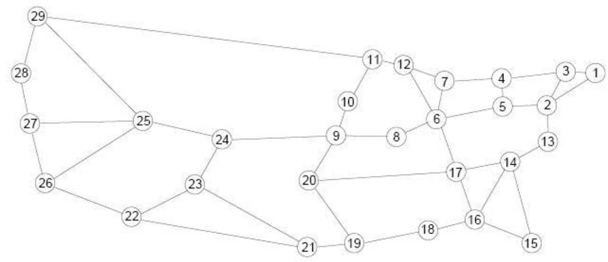

Figure 2. G29

Moreover, we give in Figures 3 to 5 the distributions of the numbers of required wavelengths obtained over the 100 runs of each method. The *x*-axis represents the number of required wavelengths *W*, and the *y*-axis represents the number of times that each value has been observed.

Let us notice that other instances have been studied (other networks and other sets of SLDs), and the same type of results have been obtained every time.

TABLE I. RESULTS FOR G57

	G57-500	G57-1000	G57-3000
Gr	44.89 - 0.028 s	67.38 - 0.083 s	91.53 - 0.307 s
RGr	42.95 - 2.68 s	65.24 - 13.54 s	88.82 - 74.54 s
Gr+	40.96 - 2.68 s	61.30 - 13.39 s	84.07 - 74.43 s
ρ(*Gr*, *Gr+*)	8.75 %	9.02 %	8.15 %
ρ(*RGr*, *Gr+*)	4.63 %	6.04 %	5.35 %

TABLE II. RESULTS FOR G29

	G29-500	G29-1000	G29-3000
Gr	43.58 - 0.018 s	65.18 - 0.056 s	78.61 - 0.198 s
RGr	41.44 - 2.16 s	62.84 - 7.65 s	76.08 - 49.85 s
Gr+	38.75 - 2.15 s	60.44 - 7.63 s	70.59 - 49.71 s
ρ(*Gr*, *Gr+*)	11.08 %	7.27 %	10.2 %
ρ(*RGr*, *Gr+*)	6.49 %	3.82 %	7.22 %

TABLE III. RESULTS FOR G200

	G200-500	G200-1000	G200-3000
Gr	19.68 - 0.031 s	50.35 - 0.170 s	80.19 - 0.887 s
RGr	18.55 - 1.68 s	48.40 - 26.76 s	77.89 - 292.30 s
Gr+	16.84 - 1.67 s	43.14 - 26.62 s	67.59 - 292.15 s
ρ(*Gr*, *Gr+*)	14.43 %	14.34 %	15.71 %
ρ(*RGr*, *Gr+*)	9.22 %	10.87 %	13.22 %

IV. CONCLUSIONS

According to the results presented above, the post-optimization method appears as improving significantly the results given by the greedy heuristic, which were already good (let us notice once again that further experiments have been carried out and lead to the same conclusions), 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 in this paper, the gain yielded by the application of *Gr+* with respect to the sole application of the greedy heuristic *Gr* (more precisely the ratio $(W_{Gr} - W_{Gr+})/W_{Gr}$) exceeds generally 10 % (the average gain is equal to 11 %), and it reaches nearly 16 % for the instance G200-3000. Even when considering the same CPU time, the gain of *Gr+* with respect to *RGr* (measured similarly by $(W_{RGr} - W_{Gr+})/W_{RGr}$) remains significant: 7.4 % in average, with a maximum of 13.2 % for G200-3000.

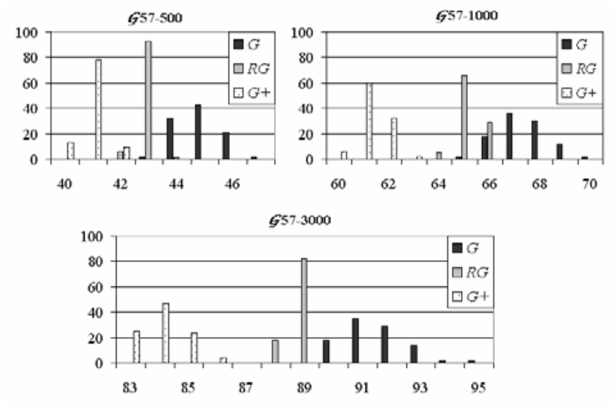

Figure 3. Distributions of the numbers of required wavelengths for G57

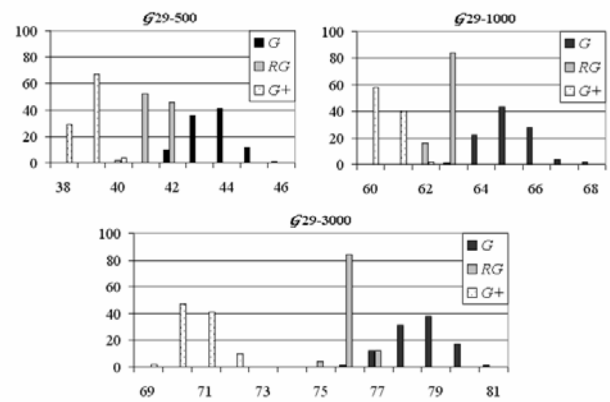

Figure 4. Distributions of the numbers of required wavelengths for G29

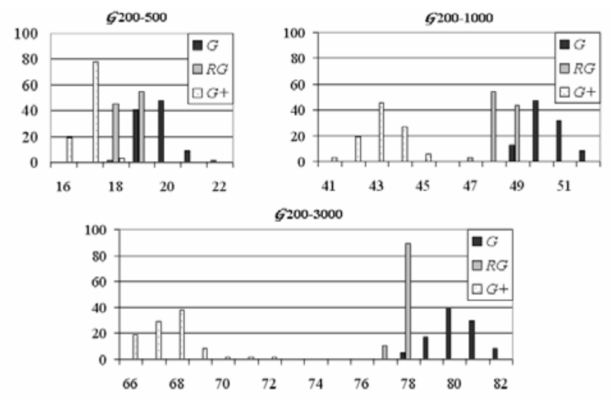

Figure 5. Distributions of the numbers of required wavelengths for G200

Another important asset of *Gr+* can be observed on the histograms (Figures 3 to 5): *Gr+* succeeds in finding some values of *W* that neither *Gr* nor *RGr* can reach during the 100 runs. When the number of SLDs increase, the gap between the histograms of *Gr+* on the one hand and those of *Gr* and *RGr+* on the other hand becomes larger. Moreover for heavy loads of traffic demands (1000 or 3000 according to the considered network), the histograms for *Gr+* become apart completely from the ones of *Gr* and *RGr*: the worst solution provided by *Gr+* remains better than the best solution found by *Gr* or *RGr*.

On the other hand, *Gr+* is significantly longer than *Gr*. In our experiments, the CPU time required to perform *Gr+* can reach few minutes, whereas it is nearly instantaneous for *Gr* (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).

However, in order to reduce the required time for the application of *Gr+*, we may modify the way to deal with the layers during the post-optimization method: instead of examining all the layers in order to empty them as much as possible, we propose to deal with only the layers corresponding to large values of wavelengths. These layers are indeed more likely to yield a gain in the total number of required wavelengths. More precisely, we introduce a parameter i varying from 0 to W; then we examine only the i layers of values $W - i + 1, ..., W - 1, W$.

Figures 6 and 7 give respectively the average number of required wavelengths and the CPU time in seconds with respect to the value of the parameter i for the instance G57-500 (results turned out to be similar for the other instances). The label N (none) on the x-axis corresponds to the result given by *Gr*, whereas the label A (all) corresponds to *Gr+* when all layers have been rearranged ($i = W$).

We observe that the CPU time varies almost linearly with respect to the number of rearranged layers i and that the quality of the solution increases when the post-optimization method is applied to more layers. However, this gain is larger for the high layers than for the low layers (the slope of the graph decreases when i increases), as expected. For the considered instance, we notice that we could have dealt with only one third (or even one quarter) of the layers, reducing thus the CPU time in the same proportion, without loosing much in terms of quality of the provided solutions. Thus we can choose the number of rearranged layers with respect to the available time.

We may conclude that the post-optimization method improves the greedy algorithm significantly and in a reasonable CPU time. Even if *Gr* is repeated, it remains clearly better to use *Gr+* than *RGr*. 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.

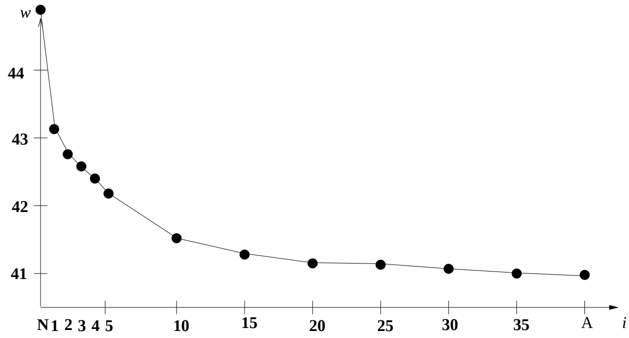

Figure 6. Results according to the number of rearranged layers for G57-500

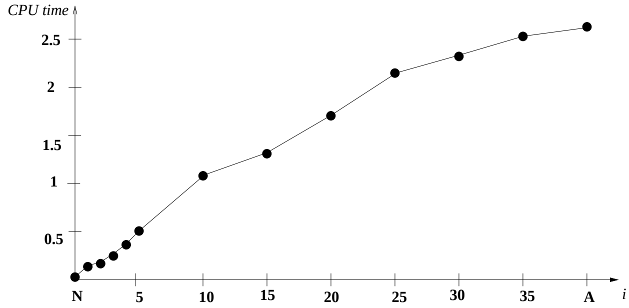

Figure 7. CPU time (in seconds) according to the number of rearranged layers for G57-500

REFERENCES

[1] D. Banerjee, and B. Mukherjee, “A practical approach for routing and wavelength assignment in large wavelength-routed optical networks,” *IEEE Journal on Selected Areas in Communications*, vol. 14 (5), pp. 903–908, 1996.

[2] L. Belgacem, 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 Geltrú, Spain, pp. 1–6, 2008.

[3] C. Chen, and S. Banerjee, “A new model for optimal routing and wavelength assignment in wavelength division multiplexed optical networks,” in INFOCOM'96, *Proceedings IEEE*, vol. 1, pp. 164-171, 1996.

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

[5] J.S. Choi, N. Golmie, F. Lapeyere, F. Mouveaux, and D. Su, “A functional classification of routing and wavelength assignment schemes in DWDM networks: static case,” *Proc. VII Int. Conf. on Optical Communication and Networks*, 2000.

[6] R. Dutta, A. E. Kamal, and G. N. Rouskas (Eds.), *Traffic Grooming for optical networks: foundations and techniques*, Springer, 2008.

[7] M. Gagnaire, J. Kuri, and M. Koubaa, *From Network Planning to Traffic Engineering in Translucent Optical WDM Networks*, Springer, 2009.

[8] H. Gi Ahn, 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*, pp. 342-351, 2005.

[9] B. Jaumard, C. Meyer, and B. Thiongane, “ILP formulations for the routing and wavelength assignment problem: symmetric systems,” in *Handbook of Optimization in Telecommunications*,

M.G.C. Resende and P.M. Pardalos Eds, Springer, 2006, pp. 637-677.

[10] R.M. Krishnaswamy, and K.N. Sivarajan, "Algorithms for routing and wavelength assignment based on solutions of LP-relaxation," *IEEE Communications Letters* vol. 5 (10), pp. 435-437, 2001.

[11] M.S. Kumar, and P.S. Kumar, "Static lightpath establishment in WDM networks - new ILP formulations and heuristic algorithms," *Computer Communications* vol. 25, pp. 109-114, 2002.

[12] J. Kuri, *Optimization problems in WDM optical transport networks with scheduled lightpath demands*, PhD thesis, ENST, Paris, 2003.

[13] J. Kuri, N. Puech, and M. Gagnaire, "Diverse routing of scheduled lightpath demands in an optical transport network," in *Proc. Design of Reliable Communications Networks (DCRN2003)*, pp. 69--76, 2003.

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

[15] K. Lee, K.C. Kang, T. Lee, and S. Park, "An optimization approach to routing and wavelength assignment in WDM all-optical mesh networks without wavelength conversion," *ETRI Journal* vol. 24 (2), pp. 131-141, 2002.

[16] L. Margara, and J. Simon, "Wavelength assignment problem on all-optical networks with k fibres per link. In Automata, Languages and Programming," *27th International Colloquium, ICALP 2000*, LNCS 1853, Springer, pp. 768-779, 2000.

[17] T. Noronha, C. Ribeiro, "Routing and wavelength assignment by partition coloring", *European Journal of Operational Reseach* vol. 171 (3), pp. 797-810, 2006.

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

[19] R. Ramaswami, and K.N. Sivarajan, "Routing and Wavelength Assignment in All-Optical Networks," *IEEE/ACM Transactions on Networking* vol. 3 (5), pp. 489--500, 1995.

[20] R. Ramaswami, and K.N. Sivarajan, *Optical Networks - A Practical Perspective*, Morgan Kaufmann, 2nd edition, 2002.

[21] C. V. Saradhi, and M. Gurusamy, "Graph Theoretic Approaches for Routing and Wavelength Assignment of Scheduled Lightpath Demands in WDM Optical Networks," *Proc. 2nd International Conference on Broadband Networks* vol. 2, 2005.

[22] N. Skorin-Kapov, *Heuristic Algorithms for Virtual Topology Design and Routing and Wavelength Assignment in WDM Networks*, hD thesis, university of Zagreb, Croatia, 2006.

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

[24] N. Skorin-Kapov, "Routing and Wavelength Assignment in Optical Networks Using Bin Packing Based Algorithms," *European Journal of Operational Research* vol. 177 (2), pp. 1167-1179, 2007.

[25] N. Skorin-Kapov, "WDM Optical Networks Planning using Greedy Algorithms," in *Greedy Algorithms*, W. Bednorz Ed., InTech Publishers, 2008, pp. 569-586.

[26] N. Skorin-Kapov, and M. Kos, "Static Routing and Wavelength Assignment in Wavelength Routed WDM Networks," in *Proceedings of the 13th IEEE Mediterranean Electrotechnical Conference (Melecon2006)*, Benalmadena (Malaga), Spain, pp. 692-695, 2006.

[27] N. Wauters, and P. Demeester, "Design of the optical path layer in multiwavelength cross-connected networks," *IEEE Journal on Selected Areas in Communications* vol. 14 (5), pp. 881-892, 1996.

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

[29] J. Zheng, and H. Mouftah, *Optical WDM Networks*, Wiley Interscience, 2004.

AUTHORS

L. Belgacem is with Futur Master, Boulogne-Billancourt, France (e-mail: Lucile.Belgacem@futurmaster.com).

I. Charon, is with the Department of Computer Science and Networks, École nationale supérieure des télécommunications, Paris, France (e-mail: Irène.Charon@enst.fr).

O. Hudry, is with the Department of Computer Science and Networks, École nationale supérieure des télécommunications, Paris, France (e-mail: IOlivier.Hudry@enst.fr).

Manuscript received 20 July 2009.

Published as submitted by the authors.