



HAL
open science

Lightpath assignment for multifibers wdm optical networks with wavelength translators

David Coudert, Hervé Rivano

► **To cite this version:**

David Coudert, Hervé Rivano. Lightpath assignment for multifibers wdm optical networks with wavelength translators. RR-4487, INRIA. 2002. inria-00072101

HAL Id: inria-00072101

<https://inria.hal.science/inria-00072101>

Submitted on 23 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

***Lightpath assignment for multifibers wdm optical
networks with wavelength translators***

David Coudert — Hervé Rivano

N° 4487 – version 2

version initiale Juin 2002 – version révisée Aout 2003

THÈME 1



***Rapport
de recherche***

Lightpath assignment for multifibers wdm optical networks with wavelength translators

David Coudert*, Hervé Rivano†

Thème 1 — Réseaux et systèmes
Projet MASCOTTE

Rapport de recherche n° 4487 – version 2‡ — version initiale Juin 2002 — version révisée
Aout 2003 15 pages

Abstract: We consider the problem of finding a lightpath assignment for a given set of communication requests on a multifiber WDM optical network with wavelength translators. Given such a network, and w the number of wavelengths available on each fiber, k the number of fiber per link and c the number of partial wavelength translation available on each node, our problem stands for deciding whether it is possible to find a w -lightpath for each request in the set such that there is no link carrying more than k lightpaths using the same wavelength nor node where more than c wavelength translations take place. Our main theoretical result is the writing of this problem as a particular instance of integral multicommodity flow, hence integrating routing and wavelength assignment in the same model. We then provide three heuristics mainly based upon randomized rounding of fractional multicommodity flow and enhancements that are three different answers to the trade-off between efficiency and tightness of approximation and discuss their practical performances on both theoretical and real-world instances.

Key-words: Multifiber optical networks, WDM, network design, wavelength translation, multicommodity flow, randomized rounding, heuristic.

* Projet commun CNRS/INRIA/UNSA MASCOTTE

† FRANCE TELECOM R&D DMI/ISE, projet MASCOTTE

‡ Affiliations des auteurs manquantes dans la version originale

Routage optique dans les réseaux wdm multifibres avec conversion partielle

Résumé : Nous considérons le problème du routage optique d'un ensemble donné de requêtes de communications dans un réseau WDM multifibres avec conversion partielle. Étant donné un tel réseau disposant de w longueurs d'onde par fibre, k fibres par lien et c conversions possibles par nœud du réseau, le problème revient à décider s'il est possible de trouver un chemin w -coloré pour chaque requête, de sorte qu'au plus w chemins utilisent une même longueur d'onde sur un même lien du réseau et qu'aucun nœud n'opère plus de c conversions. Notre résultat principal réside dans l'écriture de ce problème sous la forme d'une instance particulière d'un multiflot entier, intégrant dans un même modèle le routage et l'affectation de longueurs d'onde. Nous fournissons ensuite trois heuristiques basées sur l'arrondi aléatoire de multiflots fractionnaires, qui sont trois réponses différentes au compromis efficacité/précision des approximations. Nous les validons en comparant leur performances sur des instances théoriques et d'autres, issues du monde réel.

Mots-clés : Réseaux optiques, WDM, dimensionnement, multiflot, multifibre, conversion, arrondi aléatoire, heuristique.

1 Introduction

1.1 Wdm optical routing

Wavelength Division Multiplexing (WDM) is currently the most promising existing optical network technology, since it allows for efficient use of the high bandwidth offered by optical networks. Under WDM, laser beams carried on different wavelengths are used to implement fixed end-to-end connections — called lightpaths in this context — in the network. The major constraint imposed by this technology is that different lightpaths cannot share the same wavelength over the same fiber.

This constraint leads to mixed routing and coloring issues. Therefore the capacity of each fiber in terms of wavelengths may be dramatically under-exploited. To cope with this theoretical limitation on the efficiency of WDM optical networks, two major technologies are deployed which give more flexibility in the use of wavelength hence improving the network efficiency.

The first consists in physically interconnecting two nodes of the network with many fibers. This strategy is also motivated by the overwhelming cost of trench-digging to bury the optical fibers compared to the actual cost of a fiber. Therefore telecommunication operators rather install *multifibers networks*.

The other technique increases the complexity of nodes routers by adding *wavelength translation* equipment.

1.2 Wavelength translation

The use of a wavelength translator in a WDM network allows to change the wavelength of an incoming optical beam to another wavelength. The wavelength translation have been extensively studied on different models and has been proved that to dramatically improve the efficiency of optical bandwidth allocation [RM98, SAS99, ACG⁺01, SV00]. However, wavelength translation is a very expensive technology and it is not realistic to assume full-translation capabilities to all nodes of the network.

Therefore, two prevalent approaches are addressed in the literature concerning WDM optical networks with wavelength translation: sparse translation and limited translation. In a network with sparse wavelength translation, a fraction of the nodes are equipped with wavelength translators that are able to perform an arbitrary number of simultaneous translations [SAS99]; on the other hand a network with limited translation interconnects nodes which host wavelength translators, but these devices can only perform a limited number of translations [ACG⁺01, SV00].

1.3 Roadmap

This work addresses the optimization of the design of multifiber WDM optical networks with limited wavelength translation (MONT). Since this is a multi-criteria optimization problem, we handle it through the study of the associated decision problem of lightpath assignment

(LAP), defined in Section 2. Section 3 is dedicated to the writing of the LAP as a multicommodity flow of reasonable size compared to previous work [KS01] although we capture more generic situations. We then propose three heuristics algorithms in Section 4. These heuristics are based on *randomized rounding* of the multicommodity flow LP-relaxation, at different positions in the trade-off between computational efficiency and quality of approximation depending on how independently the variables are rounded. We then validate these heuristics in Section 5 by comparing them and the exact ILP on two theoretical and one large real-world example.

2 Multifiber lightpath assignment with wavelength translation

Connections in a WDM network are realized by laser beams carried on one available wavelength. A beam reaching a node through a fiber can be optically and passively deflected by the cross-connector and propagates toward another node through another fiber without changing its wavelength. Alternatively the beam can be “*wavelength translated*” by any kind of optical or opto-electronical mean [YZP⁺00] before being retransmit on the egress fiber.

The set of beam realizing a connection is called a *lightpath* and a *w-lightpath* if it uses at most w wavelengths.

Definition 1

Let $G = (V, E)$ be a network, w a number of wavelengths, and $u, v \in V$ two nodes of G .

A w -lightpath of G from u to v is a sequence P of $c_P + 1$ pairs $(p_i, w_i)_{i \leq c_P}$ where p_i is a path in G such that p_0 starts at u , p_{c_P} ends at v , and $\forall i \in [1..c_P]$, p_i starts where p_{i-1} stops. $\forall i$, w_i is the wavelength carrying the beam describing p_i .

We say that P is “*wavelength translated*” (or, shortly, translated) at each node where a p_i stops ($i < c_P$). Hence c_P is exactly the number of wavelength translations done on P .

The lightpath assignment problem (LAP) can be stated as follows: given a network, a set of communication requests, and a given amount of resources (wavelengths, fibers, wavelength translators), find a lightpath for each request under the constraint that only one path can use a given wavelength on a given fiber.

Depending on the way resources are accounted, the associated decision problem is written differently.

Definition 2 (Lightpath assignment problem)

Input: A network $G = (V, E)$, a set of communication requests I , a number of wavelength w , a function $k : E \rightarrow \mathbb{N}^*$ giving the number of fibers on any link of G , and (a) a function $c : V \rightarrow \mathbb{N}$ giving the number of translations available at any node of G or (b) the maximum number of translations available in G , c_{max} .

Output: Decide if it is possible to find, for each $(x, y) \in I$, a w -lightpath from x to y such that

1. On any link $e \in E$, at most $k(e)$ lightpaths use the same wavelength:
 $\forall i, \forall e, |\{P \ni (e, i)\}| \leq k(e)$;
2. (a) For any node $u \in V$, the number of wavelength translations happening at u is at most $c(u)$;
or
(b) The total number of wavelength translations ($\sum_P c_P$) is at most c_{max} .

The LAP has been extensively studied on single-fiber networks ($\forall e, k(e) = 1$) with or without wavelength translation [RM98, SAS99, KS01, SV00], but has not in the case of multifiber networks where the *wavelength assignment problem* as been mainly addressed [MS00, LS00, FPR⁺01, Riv01].

Remark 2.1 *One should remark that a w -lightpath may follow a non-simple path in the network even in the case where the p_i paths are simple. For instance p_i can go back along a part of p_{i-1} . Surprisingly a non-simple path may be part of an optimal solution of LAP though it is an absurd object when wavelength translation is not available.*

Indeed let us suppose that we are in the (a) case of Definition 2: we are given a function $c : V \rightarrow \mathbb{N}$ upper bounding the authorized number of wavelength translations at each node. In this case, a path between u and v may go through a node x where it should be translated. If all translations at x are already used, then one can “pick-up” a translation at another node y by following a path from x to y using the original wavelength w_1 , then be translated from w_1 to w_2 at y and finally following a path from y to x with the wavelength w_2 , the resulting lightpath looking like $(u \rightarrow x \rightarrow y, w_1), (y \rightarrow x \rightarrow v, w_2)$.

Hence non-simple paths may be a valid solution when one tries to minimize the maximal complexity of a node (case (a) of Definition 2) but will not be if the objective is to minimize the total number of translations (case (b) of Definition 2). Please note that using non-simple path is not non-sense from a technological point of view.

3 Lightpath assignment and Flow

Most communication problems have plenty of different writings as integer linear programs (ILP) but most of them are naive translations of choices into binary variables, leading to programs which linear relaxations hardly give informations on integral solutions. The only interest of these ILP formulations lies in the existence of efficient solvers. The LAP could be another example of this “rule” without any better understanding of the problem.

In the following we prove that LAP is a multicommodity flow problem in an auxiliary graph. Our work is inspired by the study of the routing of a multicast¹ on single-fiber

¹It could be more accurate to speak about *multi-unicast* as far as we want to establish a path between the source and each recipient.

translation-free WDM networks presented in [BHP98]. Recently, a very similar idea has been applied to general communication pattern on the same networks by [KS01].

We first generalize the case of multicast patterns to general MONT where we solve the LAP with a flow. Afterward we extend this to the general case where we use a multicommodity flow which is slightly more efficient than the one in [KS01] despite the generality.

3.1 Multicast and general requests

At the first glance, the basic part of a communication pattern is the request between two nodes which is affected a lightpath. Therefore there are in general a quadratic number of different kinds of requests in the network, one per pair of nodes. On the other hand, one may aggregate this basic parts in order to manipulate communication patterns at a higher level, hence computationally more tractably.




Figure 1: A multicast request.

Indeed, a communication pattern I can be split into the set of multicasts $M_s = \{(s, y) \mid (s, y) \in I\}$, each of them gathering all requests of source s .

We show in the following that the LAP can be solved by a flow in an auxiliary graph when the communication pattern is a multicast.

3.2 Multicast and flow

Our writing of the LAP on multicast as a flow problem uses the following three key ideas.

1. The capacity used by n units of flow going from node s to t in a flow network can be decomposed into n simple paths from s to t using one unit of capacity.
2. A set of paths that can use the same wavelength in a WDM network has load on link e (*i.e.* the number of paths going through e) at most $k(e)$ the number of fibers on this link.
3. A set of paths given by the decomposition of a flow has a load on link e exactly the same that the capacity used by the flow on e .

Therefore, given a network G , a multicast pattern M rooted at a source node s , and all the resources of the network (w , $k(e)$ and $c(e)$), we will build a flow network such that there is a flow of $|M|$ units between the source and the sink if and only if there exists a lightpath between s and each recipient within the resource limits. Figure 2 describes the flow network that solve the LAP for the network and the multicast pattern shown in Figure 1 when 3 wavelengths are available (for purpose of readability only one “*translator widget*” is drawn. There is one for each node).




Figure 2: The associated flow network with $w = 3$.

Ideas 2) and 3) drive us to consider a flow network \mathcal{N} where there are w copies of G with capacity $k(e)$ for all copies of the link e . A beam in G using wavelength i can be modeled by a unit of flow in the i^{th} copy of G in \mathcal{N} .

Obviously a set of beams using w wavelengths and $k(e)$ fibers on any link e can be modeled by a set of units of flow in \mathcal{N} where the capacities are respected.

To cope with w -lightpaths, we add a *translator widget* for any node u of G . Such a widget is made of 2 nodes trans_u^+ and trans_u^- , a link between these two nodes with capacity $c(u)$ and $2w$ uncapacitated links from each u_i to trans_u^+ and from trans_u^- to u_i .

A lightpath which is translated from wavelength w_1 to w_2 at node u can be modeled by previously described units of flow for the beams and a unit of flow from u_{w_1} to u_{w_2} via trans_u^+ and trans_u^- .

Obviously if a set of lightpaths uses less than $c(u)$ translations at any node u , it is modeled by a set of units of flow on \mathcal{N} respecting all capacities.

In order to find lightpaths between nodes s and t in G , we add a node S , a link from S to each s_i , a node super_t , and a link from each t_i to super_t . To any lightpath from s to t we can associate a unit of flow from S to super_t and conversely (using idea 1). Hence the lemma.

Lemma 1 *Sending d units of flow between S and super_t in \mathcal{N} is equivalent to find d w-lightpaths from s to t in G .*

Considering a multicast pattern M rooted at s , \mathcal{N} is the flow network where there is a node super_t for all $t \in M$ connected as described previously and a last node S' connected to all super_t nodes by a link of capacity $d(t)$ the number of lightpaths to be found from s to t .

Theorem 2 *Given a network G , a multicast pattern M rooted at a source node s , and all the resources of the network (w , $k(e)$ and $c(e)$), solving the LAP is equivalent to find in the auxiliary network \mathcal{N} a flow from S to S' of size $\sum_M d(t)$.*

Proof : From the lightpaths to the flow is straightforward.

From the flow to the lightpaths is forced because the only links reaching S' come from super_t , $t \in M$, with capacity $d(t)$. Therefore, if there is a flow of size $\sum_M d(t)$ between S and S' , exactly $d(t)$ units of flow goes through each link from super_t to S' . Hence, for each $t \in M$ there is $d(t)$ units of flow from S to super_t . Lemma 1 gives the conclusion. \square

With few simplifying of the linear program, we can write the LAP as the following ILP.

Integer Program 1

Kirchoff laws:

$$\begin{aligned}
& \forall u \in V \setminus V_s, \forall \omega < w, \\
& \quad \sum_{\Gamma^+(u)} f_\omega(e) + C_\omega^+(u) - \sum_{\Gamma^-(u)} f_\omega(e) - C_\omega^-(u) = 0 \\
& \forall u \in M_s, \forall \omega < w, \\
& \quad \sum_{\Gamma^+(u)} f_\omega(e) + Out_\omega(u) - \sum_{\Gamma^-(u)} f_\omega(e) = 0 \\
& \quad \sum_{\Gamma^+(s)} f_\omega(e) - \sum_{\Gamma^-(s)} f_\omega(s) - In_\omega = 0 \\
& \forall u \in V \setminus V_s, \\
& \quad \sum_{\omega < w} C_\omega^+(u) - \sum_{\omega < w} C_\omega^-(u) = 0 \\
& \forall u \in M_s, \\
& \quad \sum_{\omega < w} Out_\omega(u) - d(u) = 0 \\
& \quad \sum_{\omega < w} In_\omega - \sum_{u \in M_s} d(u) = 0
\end{aligned}$$

Capacity constraints:

$$\begin{aligned}
& \forall e \in E, \forall \omega < w, f_\omega(e) \leq k(e) \\
& \forall u \in V \setminus V_s, \sum_{\omega < w} C_\omega^+(u) \leq c(u)
\end{aligned}$$

3.3 General case and multicommodity flow

We now address the case of a general communication pattern $I = \bigoplus_{s \in E} M_s$. We generalize the previous network \mathcal{N} by adding a group of nodes S , S' for each multicast M_s . The LAP is now equivalent to finding an integral flow of $\sum_{t \in M_s} d(s, t)$ units commodities between S and S' for each $s \in E$. The general shape of \mathcal{N} is shown in Figure 3.

The resulting ILP has $O(|V|. (|E| + |V|). w)$ variables and $O(|V|^2.w)$ constraints which is small compared to the $O(|V|^2. |E|. w)$ and $(|V|^3.w)$ of the one presented in [KS01] which does not capture multifibers networks nor wavelength translation.

This major improvement on the size of the program has dramatic consequences on the tractability of the computing for both the exact ILP and our *randomized rounding based* approximation algorithms. Therefore we can address larger networks and communication patterns.




Figure 3: General shape of the associated flow network.

4 Approximation

Multicommodity flow problems are \mathcal{NP} -hard and among the hardest. Nonetheless theoretical works on approximation have given strong results.

The *randomized rounding* of the multicommodity flow has been proved to achieve a good theoretical approximation ratio despite the simplicity and obvious sub-optimality of the process [MR95]. This algorithm first solve the LP-relaxation to obtain fractional flows. A fractional version of the key ideas given in Section 3.2 claims that a fractional flow of n units from s to t can be decomposed into many paths with positive fractional weights summing up to n . The randomized rounding essentially consists in selecting at random one of these paths with probability its weight divided by n . A simple analyze shows that the integral capacity obtained after randomly rounding a whole multicommodity flow is a sum of Bernoulli trials with expectancy the fractional capacity. Therefore the gap between the rounded solution and the fractional one is low with high probability.

The main computational drawback of this process is that the number of paths can be exponential. Hopefully one can prove with a martingale argument that these paths are not explicitly required. The fractional unit of flow describes in the network a weighted *directed acyclic graph* (DAG) which requires linear time computation to be generated from the flow variables. A random walk inside this DAG gives a path which could be selected by the randomized rounding. If the probability to go through an edge is proportional to its weight the path is selected with the same probability as the randomized rounding.

We have implemented this approximation algorithm which has the nice property to run fairly quickly since it only needs to solve the LP-relaxation once and then does $\sum_s \sum_{M_s} d(s, t)$ random walks of length $O(|V|.w)$. It is interesting to note that the computation of the random walks is negligible compared to the solving of the LP-relaxation on any size of network and communication pattern and even using one of the best LP solver, CPLEX.

Though this approximation algorithm is very efficient, it is obvious that it is suboptimal since all path selections are done independently. Our idea is to take previous choices into account when choosing a path, hence avoiding to over-load links that have already highly loaded. The first technique we have implemented (HEUR1) is at the opposite of the randomized rounding in the trade-off between efficiency and tightness. HEUR1 consists in modifying the capacity functions k and c of the flow network after the choice of a lightpath which has non-integral weight in the fractional solution, and to recompute the LP-relaxation of the multicommodity flow where the relevant request has been decreased.

Obviously the optimality gap of the integral solution we generate is lower than in the regular randomized rounding. Our experiments validate this claim. Unfortunately, this approximation improvement is obtained at two costs. The first is that we do not know how to obtain a tight analysis of the approximation ratio. The best we can do is obtained by a straightforward and loose upper-bounding in a martingale process and drives to the same result as the randomized rounding. The second cost is computational since we may have to solve many LPS.

We have implemented a mixed heuristic, HEUR2, which rounds independently one unit of flow per multicast before solving a new LP, resulting in a more efficient but possibly less accurate approximation.

In the following we validate our algorithms on two theoretical examples and a large real-world network.

5 Simulation results and discussions

5.1 Benchmark instances

Our simulations were run on two networks: a 10 nodes ring network and a pan-american network interconnecting 65 cities coast to coast with 75 bidirectional links. Two communication patterns on the ring were randomly generated as follows. There is a connection from node u to node v with probability $4/5$ and each connection requests from 1 to 10 lightpaths with uniform probability. The first pattern (I_1) represents 376 lightpaths, the second one (I_2) 316.

The communication pattern on the pan-american network is a real-world instance which involves 1305 lightpaths.

The computations were run on a 933MHz PIII computer with 512Mo of RAM. The ILP were solved using CPLEX.

5.2 Discussion of results obtained




Figure 4: Runs on the 10-ring

The results of our simulation on the ring and on the pan-american network are plotted in Figs. 4, 6 and 5.




Figure 5: Running ratio of ILP vs Heur2

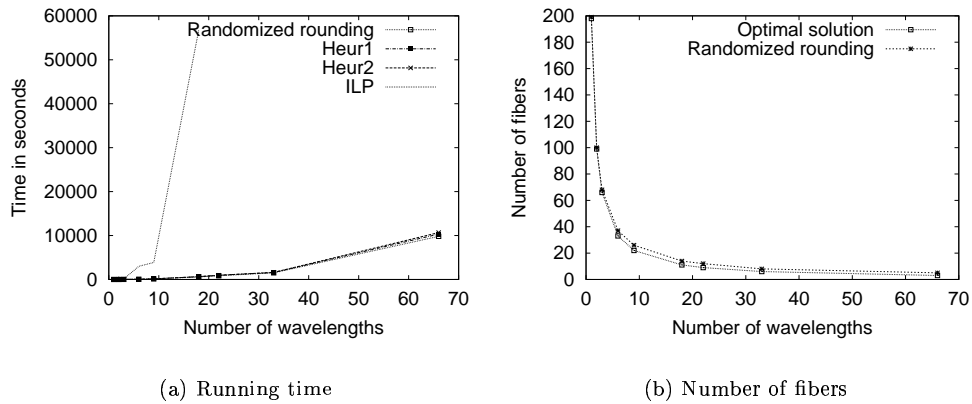


Figure 6: Runs on the pan-american network

We first remark that we were unable to find cases on the pan-american network which require translations. This is due the “simplicity” of real-world communication instances that we previously faced in other situations [FPR⁺01].

We only need translation with HEUR1 on the ring with I_2 , 1 fiber, and 56 wavelengths. In this case, the heuristic was not able to answer the decision problem without adding either 2 wavelengths or 1 translation per node.

Therefore, we plot in Figs. 4(a) and 4(b) the results of our simulations on the ring with I_1 with no translation. Fig. 4(b) depicts the minimum number of fibers needed to solve the LAP as a function of the number of available wavelengths while Fig. 4(a) gives the running time of CPLEX to solve the ILP and of our heuristics to answer the LAP for each of these cases.

Not surprisingly our heuristics are far quicker than solving the ILP. More unexpected is the dependency of the ILP solving time on the number of wavelength. The more wavelengths, the bigger the ILP, hence one could expect an ever increasing solving time as one can see in the pan-american case depicted in Fig. 6(a). This weird shape can also be observed when the communication pattern is I_2 . The time ratio between CPLEX on the ILP and HEUR2 is depicted in Fig. 5 with the number of fibers and the number of wavelengths in [4, 14]. We believe that this behavior is due to the particular structure of the network since our formulation is too general to be the most relevant here.

Nonetheless, our heuristics almost always gives the optimal solution and the randomized rounding as a gap of at most 3 fibers per link. In the case of a single-fiber translation-free ring, this fit the theoretical analysis of the similar problem of *circular arc graph coloring* made in [Kum98].

The results obtained on the pan-american network, plotted in Fig. 6, are more expected. Our heuristics are exponentially more efficient than CPLEX on the ILP and allow us to handle much larger networks and communication pattern. For instance CPLEX was not able to solve the ILP with more than 22 available wavelengths while our heuristics could handle 66. Note that the limit was due to huge computation time as well as overwhelming memory requirement.

We believe that we can go much further with the *edge-path* writing of the multicommodity flow LP and efficient operational research techniques like column generation.

Acknowledgment

We would like to thanks Stéphanne Pérennes and Afonso Ferreira for their always fruitful discussions.

References

- [ACG⁺01] 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.
- [BHP98] B. Beauquier, P. Hell, and S. Perennes. Optimal wavelength-routed multicasting. *DAMATH: Discrete Applied Mathematics and Combinatorial Operations Research and Computer Science*, 84, 1998.
- [FPR⁺01] A. Ferreira, S. Pérennes, A. Richa, H. Rivano, and N. Stier. On the design of multifiber wdm networks, August 2001. INRIA research report 4244.
- [KS01] R. Krishnaswamy and K. N. Sivarajan. Algorithms for Routing and Wavelength Assignment Based on Solutions of LP-Relaxations. In *IEEE Communications Letters*, volume 5, pages 435–437. IEEE, October 2001.
- [Kum98] V. Kumar. Approximating arc circular colouring and bandwidth allocation in all-optical ring networks. In *First International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'98)*, 1998.
- [LS00] G. Li and R. Simha. On the Wavelength Assignment Problem in Multifiber WDM Star and Ring Networks. *IEEE Infocom*, 3:1771–1780, 2000.
- [MR95] R. Motwani and P. Raghavan. *Randomized Algorithms*. Press Syndicate of the University of Cambridge, 1995.
- [MS00] L. Margara and J. Simon. Wavelength assignment problem on all-optical networks with k fibres per link. In *ICALP'00*, pages 768–779, 2000.
- [Riv01] H. Rivano. Planification de réseaux optiques wdm k-fibres. In *AlgoTel'01 - 3^{èmes} Rencontres Françaises sur les Aspects Algorithmiques des Télécommunications*, pages 41–46, Saint Jean de Luz, France, mai 2001. INRIA.
- [RM98] B. Ramamurthy and B. Mukherjee. Wavelength conversion in WDM networking. *IEEE Journal on Selected Areas of Communication*, 16(9):1061–1073, September 1998.
- [SAS99] S. Subramaniam, M. Azizoglu, and A. Somani. On optimal converter placement in wavelength-routed networks. *ACM/IEEE Transactions on Networking*, 7(5):754–766, 1999.
- [SV00] V. Sharma and E. Varvarigos. An analysis of limited wavelength translation in regular all-optical WDM networks. *IEEE/OSA Journal of Lightwave Technology*, 18(12):1606–1619, December 2000.

- [YZP⁺00] J. Yu, X. Zheng, C. Peucheret, A. T. Clausen, H. N. Poulsen, and P. Jeppesen. 40-Gbit/s all-optical wavelength conversion based on a nonlinear optical loop mirror. *IEEE/OSA Journal of Lightwave Technology*, 18(7):1001–1006, July 2000.



Unité de recherche INRIA Sophia Antipolis

2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur

INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399