Telecommunication Systems 24:2–4, 123–138, 2003

Models, Complexity and Algorithms for the Design of Multi-fiber WDM Networks *

A. FERREIRA, S. PÉRENNES and H. RIVANO ** Herve.Rivano@sophia.inria.fr
Mascotte Project, CNRS-I3S–INRIA, France

A.W. RICHA
Department of Computer Science and Engineering, ASU, USA

N. STIER MOSES
Operations Research Center, MIT, USA

Received January 2000

Abstract. In this paper, we study multi-fiber optical networks with *wavelength division multiplexing* (WDM). We extend the definition of the well-known Wavelength Assignment Problem (WAP) to the case of k fibers per link and w wavelengths per fiber, generalization that we will call (k, w)-WAP. We develop a new model for the (k, w)-WAP based on *conflict hypergraphs*. Furthermore, we consider two natural optimization problems that arise from the (k, w)-WAP: minimizing the number of fibers k given a number of wavelengths w, on one hand, and minimizing w given k, on the other. We develop and analyze the practical performance of two methodologies based on hypergraph coloring.

Keywords: optical network design, wavelength division multiplexing

Introduction

Wavelength division multiplexing (WDM) is one of the most promising optical network technologies. This technology is capable of transmitting multiple signals through the same fiber by using different wavelengths for each of them. As the interference among signals is minimal, WDM permits an efficient use of the high bandwidth offered by optical fibers. Our work focuses on WDM network design in real-life scenarios, both from theoretical and practical perspectives. Of course, operators are interested in minimizing the costs incurred in the network configuration. This leads to a design problem commonly known as the *wavelength assignment problem* (WAP), which received considerable attention in the literature [Robertson and Seymour, 22; Beauquier et al., 2; Caragiannis et al., 3; Ramaswami and Sivarajan, 21].

* A preliminary version of this paper appears in *Proceedings of the 10th International Conference on Telecommunications*, Tahiti, French Polynesia, 2003.

** Corresponding author. Mascotte Project, CNRS-I3S-INRIA, BP 93, 06902 Sophia Antipolis Cedex, France. H. Rivano is also with France Telecom R&D.

Assume that an operator possesses a WDM network and clients who request to route traffic between pairs of nodes. First, the operator establishes paths, called *lightpaths* in this context, to route each demand's traffic. After lightpaths are determined and before the network can start transmitting, the operator has to solve the WAP which consists in assigning wavelengths to lightpaths. The assignment must be done in such a way that any two paths that meet in a link are assigned different wavelengths, otherwise their signals would scramble making decodification impossible. In this paper, we consider the problem of determining paths already solved and concentrate on the one of assigning wavelengths. It is not hard to see that under the usual assumption that fibers have unlimited capacity (i.e., they can transmit an arbitrary number of wavelengths), the WAP is equivalent to the *path coloring problem* in standard graphs [Chlamtac et al., 4]. Moreover, these problems are also equivalent to the well-known *graph coloring problem* by mapping the n lightpaths to a graph of n vertices, where two vertices are adjacent if their corresponding lightpaths meet in an arc. This reduction implies that for some number $\delta > 0$, there cannot exist an n^{δ}-approximation algorithm for the WAP unless $P = NP$ [Hochbaum, 8]. As a consequence, a large fraction of the work in this field concentrated on specific topologies like line networks, rings, trees and meshes; and on specific communication patterns [Kumar, 10; Erlebach et al., 7].

From the telecommunications operator viewpoint, one of the largest costs incurred while deploying an optical network stems from physically trench-digging to bury the optical fibers. Hence, to protect themselves from demand uncertainties and failures, operators usually install many fibers. Although frequently fibers are used independently, the opportunity to exploit this redundancy gives rise to *multi-fiber* WDM *networks* (MWNs). Unfortunately, the existing work on single-fiber networks cannot be extended to MWNs in a straightforward manner. Indeed, the model used for the WAP on single-fiber networks fails to fully capture the benefits of having more fibers per link. In fact, in contrast to single-fiber networks, the presence of multiple fibers adds an extra degree of freedom when choosing wavelengths for paths: the same wavelength may be used in each of the different fibers.

Margara and Simon [16], Li and Simha [13], and Choi et al. [5] studied some theoretical properties of the new setting. For instance, they proved that increasing the number of fibers per link often simplifies the optical routing problem. In particular, although the WAP is *NP*-complete on single-fiber undirected stars, it becomes polynomial if a second fiber is available on each link. Note that using k fibers per link immediately allows a reduction of the number of wavelengths by a factor of k. In fact, multiple fibers may reduce the number of required wavelengths even further: for all k and w, there exist a network and a set of communication requests such that exactly w wavelengths are necessary to solve the problem with k fibers per link while one wavelength is enough with $k + 1$ fibers. More generally, [Margara and Simon 17] proved that for any network $\mathcal{N}$, there exists a number $k(\mathcal{N})$ such that any set of paths in $\mathcal{N}$ admits a wavelength assignment with $k(\mathcal{N})$ fibers per link.

Unfortunately, results of this flavor, which specifically determine the impact of having multiple fibers, either hold for very specific networks (e.g., [Margara and Simon, 16;

Li and Simha, 13]) or are not useful in practice [Margara and Simon, 17]. Nevertheless, some practical results can be found in the literature too. For example, Baroni et al. [1] consider path length constrained routing, wavelength assignment, wavelength conversion, and link failure restoration. They present an integer program and heuristics to minimize the total number of fibers used in the network. In addition, Hyytiä and Virtamo [9] propose the utilization of metaheuristics as simulated annealing and tabu-search for MWNs design. Both papers show that adding fibers could improve the network efficiency. Saad and Luo [23] address the dual problem of maximizing the number of lightpaths that can be established on a given network. They use a Lagrangian decomposition of certain integer programs to obtain tractable heuristics.

Another problem previously considered was that of dynamic traffic in which lightpaths have to be set-up and released dynamically. Kuri et al. [11] addressed dynamic and deterministic traffic. In this case, set-up and release times of lightpaths are assumed to be known in advance, which allows a tabu-search metaheuristic to solve the optimization problem off-line. Zhang and Qiao [25], and Li and Somani [14] study stochastic traffic using an on-line approach. They show that multi-fiber networks are more efficient than single-fiber networks with the same capacity per link (the capacity of a multi-fiber link is the sum of the capacities of each fiber in the link). The use of multi-fiber links leads to a performance that is equivalent to the one provided by limited wavelength conversion.

Most of the methodologies that we mentioned are based on heuristics and therefore it is not possible to compute performance guarantees for the solutions that they provide. This issue cannot be addressed without introducing a more formal modeling technique. For that reason, in this paper we propose an approach that mixes theory with an operational viewpoint. We build a general framework and propose a new hypergraph for modeling wavelength conflicts that arise in MWNs. To validate these concepts, we illustrate our approach with experiments on two real-world backbone networks: the European COST 239 and the Pan-American. Indeed, we solve the corresponding optimization problems with both exact integer programming formulations and approximation algorithms with bounded factor. For the former, we consider commercial LP/IP solvers while for the latter we implement the two algorithms that we propose. The first of them is a randomized algorithm whose approximation guarantee depends on the logarithm of the routing load, defined as the maximum number of paths going through an edge. The second uses randomized rounding to get a good solution, followed by re-coloring to make it feasible. Its guarantee is the best-known approximation factor for this problem, namely the logarithm of the length of the longest path.

The organization of this paper can be summarized as follows. First, in section 1, we precisely define the (k, w)-WAP in MWNs that we previously described informally. Furthermore, we develop a model that generalizes the notion of a *conflict hypergraph* used for single-fiber networks. This model captures more accurately the lightpath interdependencies in multi-fiber networks. With it, we build a bridge between results in the literature about hypergraph coloring and the (k, w)-WAP. Then, in section 2, we analyze the computational complexity of our problem. In fact, we prove that minimizing the number of wavelengths w is *NP*-hard, even in the case where the number of fibers k is fixed in advance, answering an open question regarding the complexity of this problem. In sections 3 and 4, we analyze the theoretical and practical performance of different algorithms, some of which are exact while others are approximations based on the hypergraph coloring problem. Finally, we conclude and present further directions of research in section 5.

1. Problem formulation

In this section, we formally define the (k, w)-WAP, the conflict hypergraph, and some other related concepts. Let $\mathcal{N}$ be an instance of a MWN described by the graph $G = (N, L)$ and a set of communication paths $\mathcal{P}$ on G that join the origins of the demands with their destinations. Assume further that each link in L contains k fibers and that there are w wavelengths available. The (k, w)-WAP asks for an assignment of wavelengths to paths satisfying that no more than k paths using one link are assigned the same wavelength. Therefore, in the solution each wavelength will be assigned to at most one of the fibers of a link, as required. To formalize the constraints, we define the *conflict hypergraph* of the set of paths $\mathcal{P}$ as follows (see figure 1).

Definition 1. The *conflict hypergraph* $H = (V, E)$ of the paths $\mathcal{P}$ is the hypergraph given by a vertex $v \in V$ for each path $p \in \mathcal{P}$, and an hyperedge $e \in E$ for every link $\ell \in L$. The hyperedges of H contain the vertices corresponding to paths going through the corresponding link; i.e., link ℓ generates a hyperedge $e = \{v\colon\ \ell \in \text{path represented by } v\}$.

The four main parameters of the hypergraph H can be expressed in terms of $\mathcal{N}$ and $\mathcal{P}$:

- the *number of vertices* $n := |V| = |\mathcal{P}|$;
- the *number of hyperedges* $m := |E| = |L|$;
- the *rank* $t := \max_{\ell \in L} |\{P \in \mathcal{P}\colon\ \ell \in P\}|$, which is called the *load* of $\mathcal{P}$;
- the *maximum degree* $\Delta := \max_{v \in V} |\{e \in E\colon\ e \ni v\}|$. Note that $\Delta \leqslant \max_{p \in \mathcal{P}} \text{length}(p)$, which is called the *diameter* of the routing.

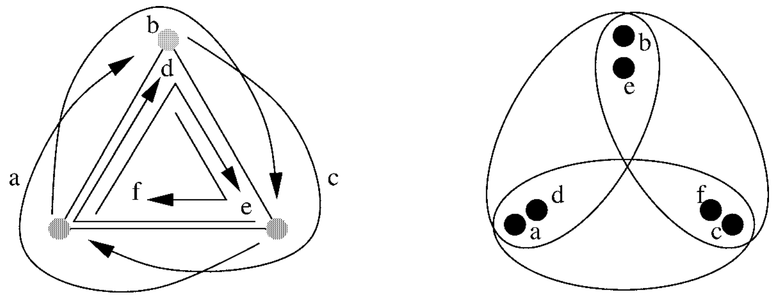

Figure 1. A ring network and the corresponding conflict hypergraph.

Notice that a vertex coloring of the conflict hypergraph induces a feasible wavelength assignment to the paths if and only if no hyperedge contains more than k vertices with the same color. This motivates the following definition.

Definition 2. Given a hypergraph $H = (V, E)$ and a set of colors $\mathcal{C} = \{1, \dots, c\}$, a mapping $f : V \to \mathcal{C}$ is a (k, c)*-coloring* if no hyperedge contains more than k vertices with the same color; i.e., for all $e \in E$ and all $\chi \in \mathcal{C}$, $|\{v \in e\colon f(v) = \chi\}| \leqslant k$.

It is easy to see from definitions 1 and 2, that there is a one-to-one correspondence between the (k, c)-colorings of the conflict hypergraph of $\mathcal{P}$ and the feasible wavelength assignments to these paths. Since we can build the conflict hypergraph H in polynomial time, this implies a polynomial time reduction from the (k, w)-WAP to the (k, c)-coloring problem. It is not trivial, however, that the converse is also true, since the hypergraph coloring problem may seem at first to be a more general (and harder) problem than the WAP. In fact, the two are equivalent and because of that we use the terms 'color' and 'wavelength' interchangeably.

2. Complexity of wavelength assignment in MWNs

In this section we prove the equivalence between the (k, w)-WAP and the (k, c)-coloring problem. Furthermore, we prove that the decision version of the (k, w)-WAP is *NP*-complete even in the case where k is fixed and present a lower bound on the number of colors needed in a (k, c)-coloring of a (hyper)clique.

Theorem 1. The (k, c)-coloring problem is polynomially equivalent to the (k, w)-WAP on MWNs.

Proof. Since we have argued in section 1 that there is a reduction from the (k, w)-WAP to the (k, c)-coloring problem, the theorem would be proven if we show the converse. Thus, we need to show that any hypergraph H is the conflict hypergraph of a set of paths $\mathcal{P}$ on a network $\mathcal{N}$, where the cardinalities of $\mathcal{P}$ and $\mathcal{N}$ are polynomial in the size of H. For instance, let $H = (V, E)$ be an arbitrary hypergraph, where $V = \{v_1, \dots, v_n\}$ and $E = \{e_1, \dots, e_m\}$. We build a network $\mathcal{N}$ and a set of paths $\mathcal{P}$ of cardinality n as follows (refer to figure 2 for an illustration).

- We create a first layer just containing a column of nodes $x_1^0, \dots, x_n^0$ from where paths will originate.
- For each hyperedge $e_i \in E$, we add one layer in $\mathcal{N}$ that contains:
  - nodes y_i and y_i' joined with an edge;
  - a column of n nodes $x_1^i, \dots, x_n^i$;
  - for each vertex $v_j \in e_i$, we add the edges (x_j^{i-1}, y_i) and (y_i', x_j^i);
  - for each vertex $v_j \notin e_i$, we add the edge (x_j^{i-1}, x_j^i).

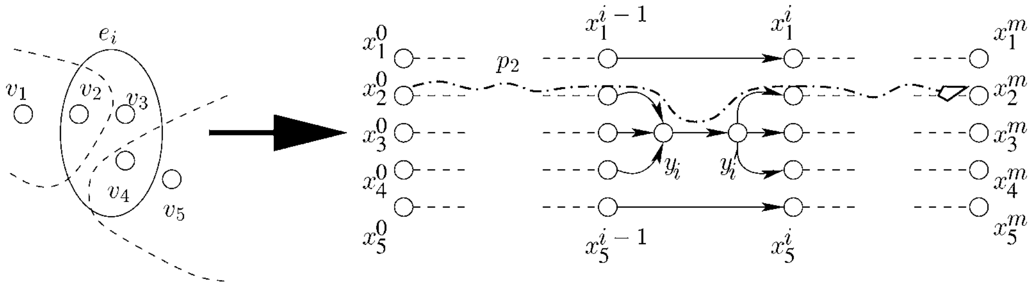

Figure 2. Reduction from a hypergraph H to a network $\mathcal{N}$ discussed in theorem 1.

- For each vertex $v_j \in V$, we add a path p_j from x_j^0 to x_j^m through nodes $x_j^1, \ldots, x_j^{n-1}$ and also through nodes y_i and y_i' when needed.

From construction, a conflict between two paths p_j and p_k can occur only on an edge of the form (y_i, y_i') such that both v_j and v_k belong to e_i. Therefore, H is exactly the conflict hypergraph of $\mathcal{P}$ on $\mathcal{N}$. Last, polynomiality follows because $|\mathcal{P}| = n$ and $|\mathcal{N}| = \mathrm{O}(mn)$. □

The graph coloring decision problem, which asks for an answer to the question "Can a graph G be colored using c colors or less?" is known to be *NP*-complete. Noting that the decision version of the (k, c)-coloring problem for $k = 1$ is exactly the coloring problem, we get the following corollary.

Corollary 2. The decision version of the (k, w)-WAP on a MWNs is *NP*-complete for an arbitrary k.

Moreover, and not too surprisingly, the problem remains difficult even with a fixed k. Indeed, the following theorem implies that the decision version of the (k, w)-WAP on a MWNs is *NP*-complete for any fixed k.

Theorem 3. The decision version of the (k, c)-coloring problem is *NP*-complete for any fixed k.

Proof. To prove the claim we reduce the (k, c)-coloring problem to the graph coloring problem. Indeed, assume we are given a graph G with n nodes and m edges. We must prove that it is c-colorable if and only if a certain hypergraph H that will define from G is (k, c)-colorable. Note that we can assume that $c < n$ because otherwise the graphs can always be colored. We will define H in such a way that it has a single (k, c)-coloring (up to permutations of the colors) and each of the c colors repeated on exactly $k + 1$ nodes. We now explain how to construct the mentioned hypergraph H.

First, we construct a hypergraph Q for which we can guarantee that there are c nodes that span all the colors. Next, we extend Q to H in a way that it is colorable if and only if G is colorable.

- Let $K_{n,t}$ denote a hypergraph with n nodes that contains all possible hyperedges of rank t. A valid (k, c)-coloring of the clique $K_{ck,k+1}$ is any coloring that satisfies that each color χ is assigned to exactly k nodes, where $\chi = 1, \ldots, c$.
- Consider the hypergraph $Q = K_{ck,k+1} \cup \{v_1, \ldots, v_c\}$; fix a valid coloring for the clique as described before; and color v_i with color i, for $i = 1, \ldots, c$. Last, add to Q all the possible $(k+1)$-hyperedges that would not contain $k+1$ nodes of the same color. By construction, it is clear that this coloring is the only possible one (up to permutations), so every color is used $k+1$ times.
- Let $H = Q$ and add to H the node-set of G. Then, for each edge (i, j) of G, we add c hyperedges as follows: for each color χ include any $k-1$ nodes that have color χ in Q and the nodes corresponding to i and j in H.

For each edge $e \in G$, the set of c hyperedges in H constraints the (k, c)-coloring problem to assign two different colors to the endpoints of e, which is exactly the constraint of the graph coloring problem. Then, if we can (k, c)-color the hypergraph, no two nodes coming from adjacent vertices of G are assigned the same color. This means that a coloring of H is valid for G and vice versa.

What remains to be seen is that the transformation is polynomial in the parameters. The total number of nodes is $c(k+1)+n$ which is polynomial in the input. Notice also that all the hyperedges we used have rank $k+1$, therefore a (loose) upper bound for the number of edges is that of the clique $K_{c(k+1)+n,k+1}$, which amounts to

$$\binom{c(k+1)+n}{k+1} \leqslant \binom{n(k+2)}{k+1} \leqslant \frac{n^{k+1}(k+2)^{k+1}}{(k+1)!}.$$

For the first inequality we used the assumption that $c < n$. As this is certainly polynomial in n and k is fixed, the claim follows. □

2.1. *A lower bound*

Extending the notion of cliques in graphs, we can give a lower bound for the number of colors needed in a (k, c)-coloring, by using (hyper)cliques, as follows.

Lemma 4. A (k, c)-coloring of $K_{n,t}$ is feasible if and only if

$$c \geqslant \begin{cases} \left\lceil \dfrac{n}{k} \right\rceil & \text{if } t > k, \\ 1 & \text{otherwise.} \end{cases}$$

Proof. The case in which $t \leqslant k$ is straightforward since when the hyperedges have rank less than k, there is no restriction on the coloring and therefore one color is enough.

For the case in which $t > k$, suppose that $K_{n,t}$ can be colored with c colors. Then, there exists a color χ such that there are $\lceil n/c \rceil$ nodes colored with it. Since the (k, c)-coloring was feasible and all hyperedges of rank t are present, we must have that $\lceil n/c \rceil \leqslant t$. Moreover, there is a hyperedge that contains $\lceil n/c \rceil$ vertices colored with χ, from where $\lceil n/c \rceil \leqslant k$. The claim follows because the minimum number that satisfies the previous condition is $\lceil n/k \rceil$.

For instance, due to the symmetry of the hypergraph, the only possible (k, c)-coloring (up to permutation of the colors) is given by assigning colors to vertices uniformly. □

The lemma above bounds the number of colors required to color any hypergraph that contains $K_{n,t}$, yielding the following generalization of the fact that the chromatic number of a graph is larger than the size of its maximum clique (just make $t = 2$ and $k = 1$).

Corollary 5. Let H be a hypergraph containing $K_{n,t}$. If H can be (k, c)-colored, with $k < t$, then $c \geqslant \lceil n/k \rceil$.

3. Tools for designing MWNs

In this section, we will present two scenarios in the design of multi-fiber networks. The number of fibers per link and the number of available wavelengths are two concurrent optimization parameters and we have no informations on how to compare their costs. Therefore we consider the two mono-criteria optimization problems where one of these parameters is fixed. The equivalence between solving the WAP for $\mathcal{P}$ and computing (k, c)-colorings of H allows us to concentrate on the latter. Hence we consider the problems of finding the minimum k (respectively c) such that there is a feasible (k, c)-coloring of H with c (respectively k) given. We address these two problems in sections 3.1 and 3.2, respectively.

3.1. Minimizing the number of fibers

We consider first the problem of minimizing the number of fibers when the number of colors is given. This optimization is related to the situation an operator faces when setting up a new network or when the operator tries to keep a set of free fibers for robustness concerns or for renting them to other clients.

Using the hypergraph model, we can formulate the problem as a minimax integer program [Srinivasan, 24]. For that, we define (0, 1)-integer variables x_{ij}, for all $i \in V$ and $1 \leqslant j \leqslant c$, such that $x_{ij} = 1$ if and only if node i is colored with color j and $x_{ij} = 0$ otherwise. The variable k is a common upper bound for the constraints defined by each hyperedge. The optimal number of fibers can be found by solving IP 1.

Integer program 1.

$$\begin{array}{lll} \text{minimize} & k & \text{(minimize \# of fibers)} \\ \text{subject to} & \displaystyle\sum_c x_{ic} = 1 & \forall \text{ node } i, \end{array}$$

$$\sum_{i \in H} x_{ic} \leqslant k \quad \forall \text{ color } c, \ \forall \text{ hyperedge } H,$$
$$x_{ic} \in \{0, 1\} \quad \forall \text{ color } c, \ \forall \text{ node } i.$$

Srinivasan [24] showed that a randomized rounding of the optimal solution of the LP relaxation of IP 1 produces, with positive probability, a solution that is feasible and approximates the optimal up to a factor of the order of the logarithm of the hypergraph degree. As said before, the degree of the conflict hypergraph of a set of paths equals to the diameter of the corresponding routing. This result uses an extension of the Lovász local lemma given in the same article. Inspired by this existential result, Lu [15] proposed a combinatorial and randomized algorithm which computes a solution within an approximation ratio not too far from the one given by [Srinivasan, 24]. This algorithm proceeds by successive recoloring phases. A first phase colors the vertices of the hypergraph randomly with a third of the set of colors. It then detects the hyperedges where there are too many vertices with the same color and colors them once again randomly with another third of the colors. Lu [15] shows then that with high probability the set of vertices that are still to badly colored if small enough for an optimal, hence exponential, coloring to be done with the last third of colors. The analysis of the algorithm shows that the approximation ratio is of the order of the logarithm of the load of the routing.

Leighton et al. [12] proposed recently an enhancement of this algorithm. Their idea comes from the remark that Lu's algorithm behaves as if it was doing a randomized rounding of a straightforward solution of the LP relaxation of IP 1 but considering the colors by thirds of the set. This strategy leads to an inefficient use of the available colors and to solutions that are worse than what promises the theoretical analysis of [Srinivasan 24]. Leighton et al. [12] want to build a solution close to the one Srinivasan showed the existence and start their algorithm with a randomized rounding of the LP relaxation. In order to be sure to get a good solution, they use recoloring techniques inspired by those of [Lu, 15]. Using linear programming for the randomized rounding is costly in terms of computing effort, but provides an efficient guide for constructing solutions within the theoretical approximation ratio. Furthermore, we shall show in section 4 that the recoloring phases are not used in practical situations since the randomized rounding yields solutions that are close to optimality. More generally, the analysis of algorithms based on randomized rounding are often very pessimistic in practice. This remark is certainly to be related to the fact that the algorithms of both [Lu, 15] and [Leighton et al., 12] can be derandomized using the conditional probability with pessimistic estimator methodology.

3.2. Minimizing the number of wavelengths

Given the number k of fibers on each link of the network, we now would like to minimize the number of wavelength required for the WAP. As an example, this optimization might model the interest of an operator who wants to upgrade the node equipments of an existing network. In terms of hypergraph coloring, this means minimizing the number of colors c such that a valid (k, c)-coloring of the hypergraph exists. We present an IP formulation for this problem below. We define a variable x_{ij} for each node i and each color j: $x_{ij} = 1$ if node i is colored with color j, and 0 otherwise. We have seen that the number of colors is bounded by $\lceil n/k \rceil$ if there are n nodes, this bound being tight when the graph is an hyperclique. We also define a variable y_j for each color j, $y_j = 1$ if there is at least one node with color j. This variables allows for counting the number of colors really used.

Integer program 2.

$$\begin{array}{lll} \text{minimize} & \sum_c y_c & \text{(minimize \# of colors)} \\ \text{subject to} & \sum_c x_{ic} = 1 & \forall \text{ node } i, \\ & \sum_{i \in H} x_{ic} \leqslant k & \forall \text{ color } c,\ \forall \text{ hyperedge } H, \\ & x_{ic} \leqslant y_c & \forall \text{ color } c,\ \forall \text{ node } i, \\ & x_{ic}, y_c \in \{0, 1\} & \forall \text{ color } c,\ \forall \text{ node } i. \end{array}$$

One could remark that IP 2 is not of high practical interest. Indeed, the program is big with $O(n^2)$ variables and $O(n^2 m)$ constraints (we could reduce the number of constraints to $O(nm)$ if the solver generates cuts automatically). Moreover, the symmetry of the formulation is critical: Mehrotra and Trick [19] have shown that the branch-and-bound techniques will waste a lot of time iterating trough similar solutions. The problem arises because after a variable is constrained by the algorithm, a permutation of the set of variables may still give a feasible solution of the same cost and this situation is uneasy to detect because of the y_j variables. Nevertheless, this issue can be addressed by automatic pruning techniques, as described by Margot [18].

4. Implementation and performance evaluation

In order to validate the analysis of the presented algorithms and the relevance of our model, we have implemented the two integer programs described in section 3, and the randomized algorithms of Lu [15] and Leighton et al., 12. This allowed us to evaluate the tradeoff between the performance and the running time of the exact version and the two approximations. We also report our findings in the experience of solving the problem of minimizing c. We ran tests on several instances and illustrate our results on two existing continental backbone networks.

4.1. The instances

The first network we report on is depicted in figure 3. COST 239 network interconnects 11 European capitals using 24 multi-fiber links. The demand matrix, which was provided

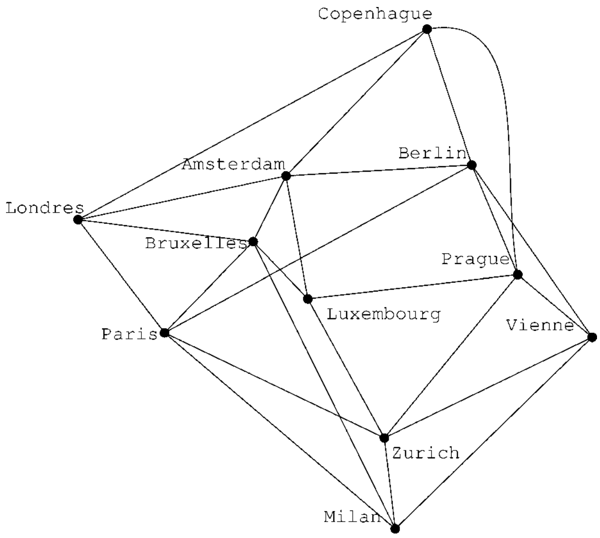

Figure 3. The COST 239 Pan-European network.

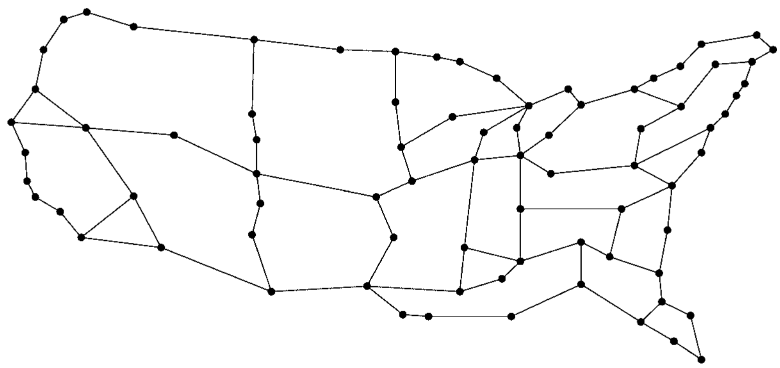

Figure 4. The Pan-American network.

to us by France Telecom [20], is made of 176 requests, which cover all possible pairs of cities. The maximum load of the given routing is 58, which is also a lower bound for the number of colors in the single-fiber case.

The second network, depicted in figure 4, is a Pan-American backbone which interconnects 78 cities with 102 links. The demand matrix was generated with the well-known gravitational model, where the weights of the cities represent their importance and are proportional to the distance to 5 main population areas in the USA. Finally, the demand between every two cities is proportional to the product of the two weights while keeping the outgoing number of requests from every city equal to the weight. Using different weights, we generated instances that were used for the benchmarks. We took reliability issues into account, and computed the routing strategy through a *minimum cost disjoint paths problem* for each origin-destination pair. For each origin and destination, the demand was randomly distributed among the disjoint paths with the shortest total distance. We report on a relatively big instance with 2022 requests with load 520.

4.2. Results

The tests we ran have validated the theoretical analysis of the algorithms and some intuition on their practical behavior. Concerning the solving of IP 1 for minimizing the number of fibers, the solver, CPLEX v7.5 found feasible solutions reasonably fast when restricted to small instances: on the European COST 239 network or the Pan-American backbone with few available wavelengths the solver did not have difficulties in proving optimality. It was expected, though, that when the instances grew bigger, the running time was going to degrade because the underlying problem is *NP*-hard. Even if this does not seem to be an issue for instances of reasonable size generated from real-world networks, it is certain that an optimal optimization of networks with more than few tens of available wavelengths is out of reach. Therefore approximate methodology are required when addressing networks deploying D-WDM technology with several hundreds of wavelengths.

Our tests gave also some precisions on the behavior of the approximation algorithms. The algorithm of Lu [15] often produces solution within an approximation ration close to the one expected from its theoretical analysis. Moreover, the strategy of dividing the set of wavelengths in 3 packs used one after the other has practical consequences on the behavior of the algorithm. The experience shows that the algorithm behaves far better when the number of wavelength is a multiple of 3 and the number of fibers computed by the algorithm decreases stepwise when the number of wavelengths increases. Therefore, the practical approximation ratio varies within a wide range. The algorithm of Leighton et al. [12] is far more efficient and we have seen that the recoloring procedures are never used in practice. Indeed, the randomized rounding always gives solutions that are very close to the optimal.

Figure 5 shows optimal and approximate computations for COST 239. It shows the required number of fibers, as a function of the number of colors available. As said

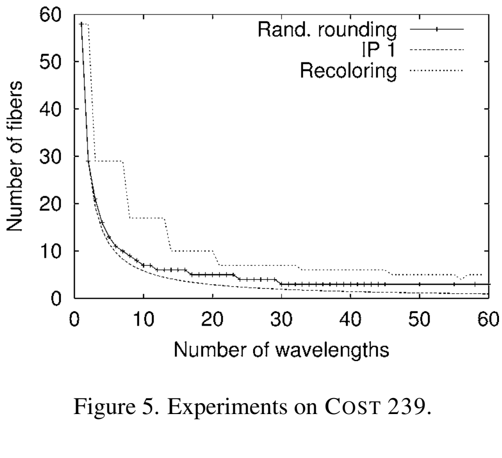

Figure 5. Experiments on COST 239.

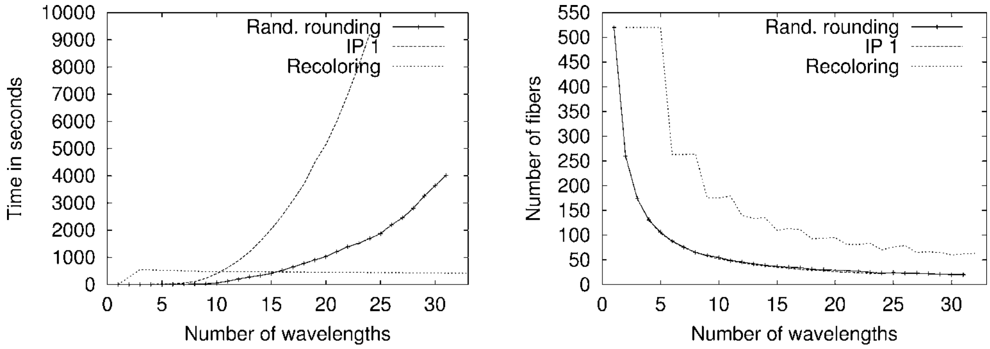

Figure 6. Experiments on the Pan-American network: time and results.

above, the small size of this network allows IP 1 to be solved to optimality even with a large number of colors. This allows for a precise comparison of the algorithms. The results furthest from optimality are those of Lu's algorithm, the stepwise dependence of the number of fibers to the number of wavelength clearly appears. The algorithm of Leighton et al. [12] produces solutions which are optimal up to an additive factor of 5. On the other hand, because of the small size of the network, both approximate algorithms finish almost immediately. Therefore, running times are not plotted for COST 239 and the running times are compared on the Pan-American network.

The big size of the instances on the Pan-American network highlights the asymptotic behaviors of the algorithms. Hence the running times of the different methodologies, depicted on the left-hand side of figure 6, clearly indicates that solving IP 1 is exponential on the number of wavelengths and is not achievable after a certain threshold, 25 wavelengths on our computers (PIV 1 GHz, 512 Mo RAM). Moreover, this impossibility is also due to the explosion of the memory space required by the branch-and-bound. The exponential growth of both the time and the memory shows that more modern computers will not solve much bigger instances and the optimization of networks with several hundreds of wavelengths has to be done approximately. The randomized rounding based algorithm of Leighton et al. [12] takes polynomial time and space which may yield difficulties with networks of high capacity, but no real impossibility. What is more surprising is that the running time of the algorithm of Lu [15] is almost constant and even slowly decreasing with the number of wavelengths. An explanation is that increasing the number of wavelengths does not increase the combinatorial complexity of this algorithm. On the contrary, the number of nodes that enter the last coloring phase may decrease with the number of wavelengths. As expected though, smaller running times are paid with the quality of approximation: the approximate number of fibers given by Lu's algorithm is around 3 times the optimal, while the algorithm of Leighton et al. [12] stays close to the optimum up to an additive factor of 4. Thus, the tradeoff between the quality of the approximation and the running time is obvious. We could use randomized rounding to optimize a static network during an offline process where the running time is not the main issue. Lu's algorithm could be useful when time is an issue, in online optimization process reacting to traffic modifications, for instance.

It is important to notice that in these instances, and often with real-world networks, the number of colors equals to its lower bound, that is, the load of the network divided by the number of fibers. It is known that pathological examples exist, but they do not usually appear in real instances.

The biggest dependency of the running time of IP 2, which optimizes the number of colors, is on the number of variables representing the colors. Initially, we used as many colors as the number of requests, because that is an upper bound. Obviously, this did not scale well when the size of the instances increased to real-world problems. Instead, we performed a binary search for the upper bound of the colors. We relied on the observation that when the bound is too small, the IP solver returns quickly that no feasible solution exists. On the other hand, when the upper bound is not tight, it takes too long to solve the first node of the branch-and-bound tree because there are too many variables. With this strategy we got IPs of the correct size that could be handled by the solver. As expected though, due to the symmetry in the formulation (the labeling of the colors can be permuted without altering the solution), the enumeration of the nodes of the branch-and-bound tree could not be always completed. In any case, we had a proof of optimality: we found that when using one less color, the LP relaxation of the problem was already not feasible. Therefore, showing a feasible solution with that many colors was enough. Indeed, it would be interesting to characterize the integrality gap of that problem.

5. Conclusion

In this paper, we have proposed a framework to model the WAP in MWNs, reducing it to a coloring problem on hypergraphs. Practically, the coloring problem appeared to be tractable when there are few colors, since its straightforward IP formulation gives optimal solutions reasonably fast. Unfortunately, this is not the case for real-world instances and the number of available wavelengths will dramatically increase with future D-WDM and UD-WDM networks.

Furthermore, the heuristics that we implemented illustrate the tradeoff between the quality of approximation and their running time. When an exact solution cannot be computed in reasonable time and space, randomized rounding can be used to produce very good solutions. When quicker solutions are required, one would rather follow approaches based on Lu's algorithm, which runs in quasi-constant time at the cost of a multiplicative factor of 3 on the optimal solution.

It is also interesting to note that these hypergraph coloring algorithms may be useful in the context of *radio ad-hoc* network optimization. Indeed, the hypergraph structure appears naturally when addressing capacity constrained radio coverage optimization. Further work is still required in this direction.

Another interesting research direction is to study the design of MWNs in the case where the routing is not fixed in advance. In such a case the lightpaths are not given, and one needs to design both the routing and the wavelength assignment at once. We believe that, as soon as k is large enough, this problem can be practically solved to optimality [Coudert and Rivano, 6].

References

[1] S. Baroni, P. Bayvel, R.J. Gibbens and S.K. Korothy, Analysis and design of resilient multifiber wavelength-routed optical transport networks, Journal of Lightwave Technology 17(5) (1999) 743–754.

[2] B. Beauquier, J.-C. Bermond, L. Gargano, P. Hell, S. Pérennes and U. Vaccar, Graph problems arising from wavelength-routing in all-optical networks, in: *2nd IEEE Workshop on Optics and Computer Science, part of IPPS'97*, Genova, Switzerland (IEEE Press, New York, 1997).

[3] I. Caragiannis, A. Ferreira, C. Kaklamanis, S. Pérennes and H. Rivano, Fractional path coloring with applications to WDM networks, in: *ICALP*, eds. F. Orejas, P.G. Spirakis and J. van Leeuwen, Lecture Notes in Computer Science, Vol. 2076 (Springer, New York, 2001) pp. 6–15.

[4] I. Chlamtac, A. Ganz and G. Karmi, Lightpath communications: An approach to high bandwidth optical WAN's, IEEE Transactions on Communications 40(7) (1992) 1171–1182.

[5] H. Choi, S. Subramaniam and H.-A. Choi, Optimal wavelength assignment algorithms for permutation traffic in multifiber WDM ring networks, Photonic Network Communications 4(1) (2002) 37–46.

[6] D. Coudert and H. Rivano, Lightpath assignment for multifibers WDM optical networks with wavelength tranlators, in: *Proc. IEEE Global Telecommunications Conf.*, Taipei, Taiwan, 2002, pp. 6–15.

[7] T. Erlebach, K. Jansen, C. Kaklamanis, M. Mihail and P. Persiano, Optimal wavelength routing on directed fiber trees, Theoretical Computer Science 221(1/2) (1999) 119–137.

[8] D.S. Hochbaum, *Approximation Algorithms for NP-Hard Problems* (PWS, Boston, MA, 1997).

[9] E. Hyytiä and J. Virtamo, Wavelength assignment in multifibre in WDM-networks, Technical Report COST257TD(99)04, Helsinki University of Technology (1999).

[10] V. Kumar, Approximating arc circular colouring and bandwidth allocation in all-optical ring networks, in: *First Internat. Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'98)*, 1998.

[11] J. Kuri, N. Puech, M. Gagnaire and E. Dotaro, Routing and wavelength assignment of scheduled lightpath demands in a WDM optical transport network, in: *Proc. of the 1st IEEE Internat. Conf. on Optical Communications and Networks*, Singapore, 2002.

[12] T. Leighton, C. Lu, S. Rao and A. Srinivasan, New algorithmic aspects of the Local Lemma with applications to routing and partitioning, SIAM Journal on Computing 31(2) (2001) 626–641.

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

[14] L. Li and A. Somani, A new analytical model for multifiber WDM networks, IEEE Journal on Selected Areas in Communications 18(10) (2000) 2138–2145.

[15] C.-J. Lu, Deterministic hypergraph coloring and its applications, in: *RANDOM*, eds. M. Luby, J.D.P. Rolim and M.J. Serna, Lecture Notes in Computer Science, Vol. 1518 (Springer, New York, 1998) pp. 35–46.

[16] L. Margara and J. Simon, Wavelength assignment problem on all-optical networks with k fibres per link, in: *ICALP*, eds. U. Montanari, J.D.P. Rolim and E. Welzl, Lecture Notes in Computer Science, Vol. 1853 (Springer, New York, 2000) pp. 768–779.

[17] L. Margara and J. Simon, Decidable properties of graphs of all-optical networks, in: *ICALP*, eds. F. Orejas, P.G. Spirakis and J. van Leeuwen, Lecture Notes in Computer Science, Vol. 2076 (Springer, New York, 2001) pp. 768–779.

[18] F. Margot, Pruning by isomorphism in branch-and-cut, in: *IPCO*, eds. K. Aardal and B. Gerards, Lecture Notes in Computer Science, Vol. 2081 (Springer, New York, 2001) pp. 304–317.

[19] A. Mehrotra and M.A. Trick, A column generation approach for graph coloring, INFORMS Journal on Computing 8(4) (1996) 344–354.

[20] Porto: Planification et Optimisation des Réseaux de Transport Optiques, Document under legal authorization of France Télécom R&D, RNRT project with the CNRS/INRIA/I3S project Mascotte, Alcatel Corporate Research Center and France Télécom R&D (2000).

[21] R. Ramaswami and K.N. Sivarajan, *Optical Networks: A Practical Perspective* (Morgan Kaufman, Los Altos, CA, 1998).

[22] N. Robertson and P.D. Seymour, Graph minors XIII. The disjoint paths problem, Journal of Combinatorial Theory Series B 63 (1995) 65–110.

[23] M. Saad and Z.-Q. Luo, A Lagrangian decomposition approach for the routing and wavelength assignment in multifiber WDM networks, in: *Proc. of IEEE Global Telecommunications Conf.*, Taipei, Taiwan, 2002.

[24] A. Srinivasan, An extension of the Lovász Local Lemma, and its applications to integer programming, in: *Proc. of the 13th Annual ACM–SIAM Symposium on Discrete Algorithms*, Atlanta, GA, 1996, pp. 6–15.

[25] X. Zhang and C. Qiao, Wavelength assignment for dynamic traffic in multi-fiber WDM networks, in: *Proc. of the 7th Internat. Conf. on Computer Communications and Networks*, Lafayette, LA, 1998, pp. 479–585.