INRIA

On the design of multifiber WDM networks

A. Ferreira — S. Pérennes — A. W. Richa — H. Rivano — N. Stier

N° 4244

August 2001

THÈME 1

Rapport de recherche

INRIA

SOPHIA ANTIPOLIS

On the design of multifiber WDM networks

A. Ferreira[*] , S. Pérennes[†] , A. W. Richa[‡] , H. Rivano[§] , N. Stier[¶]

Abstract: In this paper, we address multifiber optical networks with Wavelength Division Multiplexing (WDM). Assuming that the lightpaths use the same wavelength from source to destination, we extend the definition of the well-known Wavelength Assignment Problem (WAP), to the case where there are k fibers per link, and w wavelengths per fiber are available. We then develop a new model for the (k, w)-WAP, based on *conflict hypergraphs*: Conflict hypergraphs more accurately capture the lightpath interdependencies, generalizing the conflict graphs used for single-fiber networks. By relating the (k, w)-WAP with the hypergraph coloring problem, we prove that the former is $\mathcal{NP}$-complete, and present further results with respect to the complexity of that problem. Finally, we analyze the practical performances of two methodologies based on hypergraph coloring, on existing backbone networks in Europe and in the USA. The first relies on an integer programming formulation and the second consists of a heuristic based on a randomized algorithm. We consider the two natural optimization problems that arise from the (k, w)-WAP: the problem of minimizing k given w, and that of minimizing w given k.

Key-words: optical networks, wavelength division multiplexing, network design, optimization, wavelength assignment problem, modeling, hypergraphs, hypergraph coloring, integer programming, heuristics, COST239, Pan-American network

  * Mascotte Project, CNRS, I3S & INRIA Sophia Antipolis, BP 93 F-06902 Sophia-Antipolis, France. Email:Afonso.Ferreira@sophia.inria.fr
   † Mascotte Project. Email: Stephane.Perennes@sophia.inria.fr
   ‡ Department of Computer Science and Engineering, Arizona State University, Tempe, AZ 85287–5406, USA. Email: aricha@asu.edu. This research was conducted while this author was visiting the University of Nice - Sophia Antipolis at the Mascotte Project.
   § Mascotte Project. Email: Herve.Rivano@sophia.inria.fr
   ¶ MIT Operations Research Center, Cambridge, MA 02139, USA. Email: nstier@mit.edu. This research was conducted while this author was visiting INRIA Sophia Antipolis at the Mascotte Project.

Une étude sur le dimensionnement
de réseaux WDM multifibres

Résumé : Nous étudions dans ce rapport de recherche les réseaux optiques multifibres utilisant la technique de multiplexage en longueurs d'onde (WDM pour *Wavelength Division Multiplexing*). Sous l'hypothèse que les chemins optiques utilisent la même longueur d'onde de la source à la destination (i.e. nous ne nous préoccupons pas de question de conversion de longueur d'onde), nous étendons la définition classique du problème de planification optique (WAP pour *Wavelength Assignment Problem*) au cas où k fibres sont présentes dans chaque câble du réseau, et où w longueurs d'ondes sont disponibles sur chaque fibre. Nous développons alors une nouvelle modélisation du (k, w)-WAP basée sur *l'hypergraphe des conflits*: l'hypergraphe des conflits est une généralisation du graphe des conflits utilisé dans le cadre des réseaux monofibres qui rend mieux compte des interdépendances entre les chemins optiques. En reliant le (k, w)-WAP à un problème de coloriage d'hypergraphe, nous prouvons que le premier est NP-complet et nous présentons d'autres résultats de complexité à propos de ce problème. Ensuite nous analysons les performances concrètes de deux méthodologies basées sur de la coloration d'hypergraphe, sur des réseaux d'infrastructures déployés en Europe et aux États-Unis. La première méthodologie s'appuie sur une formulation en programmation linéaire en nombres entiers tandis que la seconde est une heuristiques issue d'un algorithme aléatoire. Nous considérons les deux problèmes d'optimisation découlant naturellement du (k, w)-WAP: la minimisation de k sachant w et celle de w sachant k.

Mots-clés :   réseaux optiques, multiplexage en longueurs d'onde, dimensionnement de réseau, optimisation, problème d'affectation de longueur d'onde, problème de planification optique, modélisation, hypergraphes, coloration d'hypergraphes, programmation linéaire en nombres entiers, heuristiques, COST239, réseau Pan-American

1 Introduction

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, 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 link.

In order to avoid wavelength conflicts, two main techniques may be used. The first technique places wavelength converters at carefully chosen network nodes, in order to ensure that all lightpaths going through a link have different wavelengths. Wavelength conversion has been the subject of a large volume of recent research (e.g., [KK01, RM88]). In practice, the major drawback of this technology lies on the very high costs of acquiring wavelength converters.

The second technique avoids the use of wavelength converters, ensuring that the lightpaths use the same wavelength from source to destination. This scenario leads to the design problem known as the *wavelength assignment problem* (WAP) [RS95, BBG$^+$97, Kum98, BHP98, RS98, Bea99, ACKP00, CFK$^+$01]. The off-line version of the WAP can be define as follows: Given a WDM network $\mathcal{N}$ and a set of lightpaths satisfying traffic requests, assign wavelengths to the lightpaths so that any two paths that cross the same link are assigned different wavelengths.

Our work focuses on studying WDM networks in real-life scenarios [PCR00], from both theoretical [BHP98, CFK$^+$01] and practical [PCR00, BFM98, CFM00b, CFM00a] perspectives. Perhaps surprisingly, 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, it is usual to have many fibers deployed between any two points of the network, giving rise to *multifiber* WDM *networks* (or MWNs for short).

Unfortunately, the existing work on single-fiber network design cannot be extended to MWNs in a straightforward way. Lately, the real impact of MWNs has attracted considerable attention [ZQ98, LS99, BBGK99, HV99, MS00, LS00, MS01]. Note that using k fibers per link immediately allows for reducing the number of wavelengths by a factor of k. Actually multifibers may allow us to reduce the number of wavelengths required even further. For example, adding just one fiber to a single-fiber network can decrease the number of wavelengths required to route n lightpaths from n to 1 [MS00, LS00].

Unfortunately, results of this flavor, which specifically determine the impact of having multifibers either hold for very specific networks, as the results in [MS00, LS00], or are very preliminary as far as modeling is concerned [ZQ98, LS99, BBGK99, HV99]. For instance, the model used for the WAP on single-fiber networks fails to fully capture the benefits of having more fibers per link when minimizing the total number of wavelengths used in the network in MWNs.

In this paper, we generalize the WAP to the case where there are k fibers per link, and w wavelengths per fiber are available — this generalization is called the (k, w)-WAP. Two optimization problems naturally arise from the (k, w)-WAP: the problem of minimizing the number of wavelengths used, given k, and that of minimizing the number of fibers k if we are

given w. In order to build a general framework around the (k, w)-WAP, we propose a new tool for modeling conflicts arising in wavelength utilization in MWNs, based on hypergraphs. The *conflict hypergraph*, formally defined in Section 3, is a generalization of the popular conflict graph, used for the WAP on single-fiber networks. We validate the concepts proposed in this work by considering both optimization problems (of minimizing k and of minimizing w) in two backbone networks: the European COST 239 [PCR00] and the Pan-American [BCF⁺00]. The contributions of this work can be summarized as follows:

- We formally define the (k, w)-WAP for MWNs, where either the number of fibers, or the number of wavelengths per fiber can be optimized.

- We develop a new model for the (k, w)-WAP, based on hypergraphs. The *conflict hypergraphs* will be used for modeling the interdependence of the lightpath intersections. The conflict hypergraphs more accurately capture the lightpath interdependencies, generalizing the concept of conflict graphs used for single-fiber networks.

- Using this new hypergraph model, we build a bridge between coloring results for hypergraphs in the literature and the (k, w)-WAP.

- We show complexity results for the (k, w)-WAP in MWNs. In fact, minimizing the number of wavelengths is $\mathcal{NP}$-complete, even in the case where the number of fibers is fixed in advance.

- We analyze the practical performances of two methodologies based on hypergraph coloring on existing backbone networks in Europe and in the USA. The first relies on an integer programming formulation and the second consists on a heuristic based on a randomized approximation algorithm. We analyze the feasibility of solving real-world (k, w)-WAP with existing LP/IP solvers of the market. The field is still open to new heuristics for hypergraph coloring.

In the literature, the *efficiency* of a k-fiber network is measured in terms of k and the number of wavelengths w_k required by the (k, w)-WAP, for fixed k. For a set of lightpaths with load $\mathcal{L}$, w_k lies between $\frac{\mathcal{L}}{k}$ and $\frac{w_1}{k}$. The *efficiency* of the network is then defined as $\frac{\mathcal{L}/k}{w_k}$.

The remainder of this paper is organized as follows. First, we present an overview of related work in Section 2. In Section 3, we give the problem formulation and the proposed hypergraph model. Then, in Section 4, we relate the (k, w)-WAP with hypergraph coloring, proving that the (k, w)-WAP is $\mathcal{NP}$-complete, and presenting other results with respect to the complexity of the problems. In Section 5, we address the actual problem of designing a multifiber network, with respect to the optimization of either parameter. Section 6 discusses our prototypes and their performance evaluation. Finally, we conclude and present some future work in Section 7.

2 Related work

Motivated by the very large costs of deploying WDM networks, a large volume of research has targeted design issues on these networks in the past.

In single-fiber networks, it is usual to assume that two nodes are connected by one fiber of unlimited capacity (i.e. able to carry any number of wavelengths). Hence the $(1, w)$-WAP (formerly known simply as WAP) is exactly the *path coloring problem* in standard graphs, which has been proven equivalent to the general vertex coloring problem [CGK92]. Thus, there exists a fixed $\delta > 0$ such that no approximation within n^δ is possible unless $\mathcal{P} = \mathcal{NP}$ [Hoc97].

Therefore, a large amount of work concentrated on specific topologies [Kum98, BHP98, Bea99, ACKP00, CFK+01, EJK+99, BGP+96]. On line networks, the efficiency is always equal to 1. For rings, a randomized $(1 + \frac{1}{e})$-approximation algorithm based on randomized rounding of multicommodity flows is proposed in [Kum98]. For trees, some of the current best approximation results are a deterministic $\frac{5}{3}$-approximation [EJK+99], a randomized 1.61-approximation for bounded degree trees based on fractional path coloring [CFK+01], and a randomized greedy $\frac{7}{5}$-approximation for binary trees [ACKP00]. Specific communication patterns have also been studied like All-to-All [Bea99, BGP+96] and multicast, which achieves an efficiency of 1 on general networks [BHP98].

The design of multifiber networks has recently been studied under different models and traffic assumptions [ZQ98, LS99, BBGK99, HV99, MS00, LS00, MS01]. For instance the $(1, w)-$WAP is $\mathcal{NP}$-complete on undirected stars but becomes polynomial with an efficiency of 1 if 2 fibers are available on each link [MS00, LS00].

Dynamic traffic — which means that lightpaths have to be established and released dynamically — has been studied in [ZQ98, LS99], where multifiber networks were shown to be more efficient than single-fiber networks with the same capacity[1] per link. Using multifiber links has also been shown to lead to performances equivalent to those provided by limited wavelength conversion.

In [BBGK99], an integer program and heuristics that solve the static problem are discussed. They consider path length constrained routing, wavelength assignment, wavelength conversion, and link failure restoration. The objective is to minimize the total number of fibers used in the network. Two meta-heuristic (simulated annealing and taboo-search) for MWNs design are proposed in [HV99]. Both papers show that adding fibers could improve the network efficiency.

MWNs have been studied from a theoretical point of view in [MS00, LS00, MS01]. For instance, it was proven that increasing the number of fibers per link often simplifies the optical routing problem: 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 1 wavelength is enough with $k + 1$ fibers.

Upper bounds on w_k were improved in [MS00, LS00], for specific topologies like stars and rings. Also, in [MS01] it was proven that for each fixed network $\mathcal{N}$, there exists a $k(\mathcal{N})$

---

[1] The capacity of a link is the product of the number of fibers on the link and the capacity of each fiber.

such that the efficiency is 1 for all sets of paths provided that there are $k(\mathcal{N})$ fibers per link. Unfortunately, $k(\mathcal{N})$ is usually very large (and therefore impractical).

3   Problem formulation

In this section, we formally define the (k, w)-WAP, the conflict hypergraph, and some other concepts related to those. Let $\mathcal{N}$ be an instance of a MWN, where the set of nodes is represented by N and the set of links by L. Assume, without loss of generality, that every link in the network contains the same number of fibers, denoted by k, and that a set of communication paths $\mathcal{P}$ is given. A solution to the (k, w)-WAP is an assignment of one out of w wavelengths to each path, satisfying the constraints that for every link in the network, no more than k paths using the link are assigned the same wavelength. In order to model these constraints, we define the *conflict hypergraph* $H(\mathcal{N}, \mathcal{P})$ of the set of paths $\mathcal{P}$ in $\mathcal{N}$, as follows:

Definition 1 *The* conflict hypergraph $H = (V, E)$ *of the paths* $\mathcal{P}$ *in* $\mathcal{N}$ *is a hypergraph such that each vertex* $v \in V$ *corresponds uniquely to a path* $p \in \mathcal{P}$, *and such that for every link* $\ell \in L$, *there exists a hyperedge in* E *containing the vertices that correspond to all the paths going through* ℓ *(and these are the only hyperedges in* E). $\qquad\square$

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* $\mathcal{C} = \{1 \ldots c\}$, *a mapping* $f : V \to \mathcal{C}$ *is a* (k, c)-coloring *if and only if no hyperedge contains more than* k *vertices with the same color, that is,* $\forall e \in E, \ \forall q \in \mathcal{C}, \ |\{v \in e : f(v) = q\}| \le k$. $\qquad\square$

The four main parameters of the hypergraph $H = (V, E)$ can be expressed in terms of $\mathcal{N}$ and $\mathcal{P}$:

- the *number of vertices* $n \equiv |V| = |\mathcal{P}|$,

- the *number of hyperedges* $m \equiv |E| = |L|$,

- the *rank* (maximum size of a hyperedge) $t \equiv \max_{\ell \in L} |\{P \in \mathcal{P} : \ell \in P\}| = $ load of $\mathcal{P}$,

- the *maximum degree* (number of hyperedges incident to a node) $\Delta \equiv \max_{v \in V} |\{e \in E : e \ni v\}|$. Note that $\Delta \le max_{p \in \mathcal{P}}$ length(p), which is equal to the diameter of the routing.

(Recall that the *load* of $\mathcal{P}$ is equal to the maximum number of paths passing through any one edge in the network, by definition.)

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. Thus, the (k, c)-coloring problem is at least as difficult as the (k, w)-WAP. Actually, we prove their equivalence in the next section. For that reason, in the sequel, we will use colors and wavelengths interchangeably.

4 Complexity of wavelength assignment in MWNs

In this section, we prove the equivalence between the (k, w)-WAP and (k, c)-coloring, prove that the (k, w)-WAP is $\mathcal{NP}$-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.
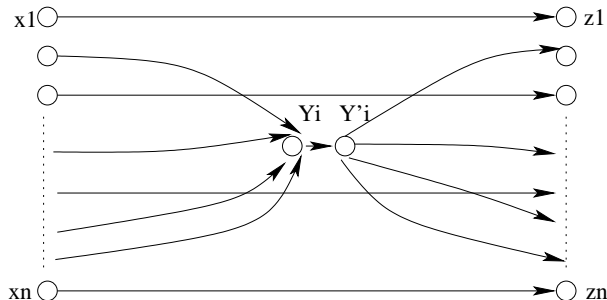
Theorem 1 *The (k, c)-coloring problem is polynomially equivalent to the (k, w)-WAP on* MWN*s.*

Proof: It is enough to prove that any hypergraph H is the conflict hypergraph of a set of paths $\mathcal{P}$ on a network $\mathcal{N}$, where the sizes of both $\mathcal{P}$ and $\mathcal{N}$ are polynomial on the size of H. Let $H = (\{v_1, \ldots, v_n\}, \{e_1, \ldots, e_m\})$ be a hypergraph, where $e_i = \{v_{j_1^i}, \ldots, v_{j_{r_i}^i}\}$, for all i. Let $r_i = |e_i|$, for all i. For every hyperedge e_i, let $\mathcal{N}_i(V_i, E_i)$ be the network (depicted in Figure 1) containing

- n nodes x_j^i, $j = 1 \ldots n$, and n nodes z_j^i, $j = 1 \ldots n$,

- two special nodes Y_i et Y_i',

- an edge $Y_i \to Y_i'$,

- $\forall j \in \{j_1^i, \ldots, j_{r_i}^i\}$,

    - an edge $x_j^i \to Y_i$, and
    - an edge $Y_i' \to z_j^i$,

- $\forall j \notin \{j_1^i, \ldots, j_{r_i}^i\}$, an edge $x_j^i \to z_j^i$.

Let $\mathcal{N}$ be the network composed of the union of the subnetworks $\mathcal{N}_i$ connected by the edges $z_j^i \to x_j^{i+1}$, $\forall i = 1 \ldots n - 1$, $\forall j = 1 \ldots n$. Clearly, the size of $\mathcal{N}$ is $O(nm)$ and H is the conflict hypergraph of the set of the unique shortest paths between x_j^1 and z_j^n, for all j, in $\mathcal{N}$. $\qquad\square$

The (k, c)-coloring problem is clearly $\mathcal{NP}$-complete for general k, since it generalizes the graph coloring decision problem when $k = 1$. Therefore,

Figure 1: *Network* $\mathcal{N}_i$

Corollary 1 *The (k, w)-WAP on a MWNs is $\mathcal{NP}$-complete for a general k.*

Moreover, we prove below that the problem remains difficult even when k is fixed.

Theorem 2 *The (k, c)-coloring problem is $\mathcal{NP}$-complete for any fixed k.*

Proof: We will reduce this problem to (standard) coloring on graphs. That is, given a graph G with n nodes and m, we have to answer "Can G be colored using c colors or less?". To prove that (k, c)-coloring is $\mathcal{NP}$-complete, we answer the graph coloring question by calling the (k, c)-coloring oracle. We can assume that $c < n$ because otherwise the answer is trivially yes.

We are given a graph G with n nodes and m edges that we must color. We extend G into a hypergraph H in the following way. Let $K_{n,t}$ be a hypergraph with n nodes that contains all the possible hyperedges of rank t. We start by adding a $K_{ck,k+1}$ clique to H, which can trivially be (k, c)-colored. Fix one of the possible colorings of the clique. We will now make that coloring the only feasible one (up to permutations of the colors). For that, we add c new nodes to H, each with a different color pre-assigned and then all the possible $(k+1)$-hyperedges that do not join $k+1$ nodes of the same color. The coloring that we fixed is, of course, feasible for this structure by construction. If we vary the cardinality of nodes having a color, that is clearly infeasible because there is a clique included. If we permute the colors there will be a hyperedge preventing that to be feasible.

The construction above allows us to claim that we have $k+1$ nodes having each of the c colors. Now, returning to the original graph, for each edge we add c hyperedges as follows: for each color χ, include any of the $k-1$ nodes that have color χ in the structure and the two endpoints of the edge. Every one of these hyperedges means that the two nodes cannot be colored using the same color, which is what we need for graph coloring.

If we can (k, c)-color the hypergraph, then we can also color the graph with c colors. What remains to be seen is that the transformation is polynomial on the parameters. We added $c(k + 1)$ nodes which is polynomial on the input. We added less than the maximum possible number of $(k + 1)$-regular hyperedges, which is equal to

$$\binom{c(k+1)}{k+1} \leq \frac{c^{k+1}(k+1)^{k+1}}{(k+1)!},$$

which is certainly polynomial on c (k is fixed). Then we added mc more hyperedges for preventing color repetitions. Therefore, recalling that $c < n$, a bound for the total number of hyperedges added is $O(n^{k+1} + mn)$, which completes the proof. $\qquad\square$

Corollary 2 *The (k, w)-*WAP *on a* MWN*s is $\mathcal{NP}$-complete for any fixed k.*

4.1 A lower bound

Extending the notion of cliques in graphs, we can give a lower bound on the number of colors needed in a (k, c)-coloring, by using (hyper)cliques, as follows. Recall that $K_{n,t}$ is a hypergraph with n nodes that contains all the possible hyperedges of rank t. The following holds.

Lemma 1 *A (k, c)-coloring of $K_{n,t}$ is feasible if and only if*

$$c \geq \left\{ \begin{array}{ll} \lceil \frac{n}{k} \rceil & \textit{if } t > k, \\ 1 & \textit{otherwise.} \end{array} \right.$$

Proof: The case where $t \leq k$ is trivial because the hyperedges do not impose any restriction on the coloring and therefore one color is enough. For the case where $t > k$, suppose $K_{n,t}$ can be colored with c colors. Then, as the hypergraph is symmetric, every color will be repeated $\lceil \frac{n}{c} \rceil$ or $\lfloor \frac{n}{c} \rfloor$ times. As we assumed that c colors were feasible, we must have that $\lceil \frac{n}{c} \rceil \leq k$ because otherwise there is a contradiction. The minimum number of colors that satisfies that condition is $\lceil \frac{n}{k} \rceil$ and it is easy to see that any permutation of the colors is feasible. $\qquad\square$

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

5   Tools for designing MWNs

In this section, we will present two scenarios in the design of multifiber networks. The equivalence between solving the WAP for $\mathcal{P}$ and computing (k, c)-colorings of H allows us to concentrate on the latter. For instance, the problems we consider are 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 5.1 and 5.2, respectively.

5.1   Minimizing the number of fibers

We consider first the problem of minimizing the number of fibers when the number of colors is given. This problem can be formulated as a Minimax Integer Program (MIP) [Sri96]. For instance, we define $(0, 1)$-integer variables x_{ij}, for all $i \in V$ and $1 \le j \le c$, 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. It is called a MIP because our goal is to minimize it. The optimal number of fibers can be found by solving the following IP.

Integer Program 1

$$
\begin{aligned}
\text{minimize} \quad & k & & \text{(minimize \# of fibers)} \\
\text{s.t.} \quad & \sum_{c} x_{ic} = 1 & & \forall \text{ node } i \\
& \sum_{i \in H} x_{ic} \le k & & \forall \text{ color } c, \ \forall \text{ hyperedge } H \\
& k \ge 0, \ x_{ic} \in \{0, 1\} & & \forall \text{ color } c, \ \forall \text{ node } i.
\end{aligned}
$$

Recently, Srinivasan showed that if the optimal solution of the LP relaxation is rounded randomly, with positive probability, a solution that is feasible and not too large can be encountered [Sri96]. A simple algorithm, discussed in Lu [Lu98], can compute a solution that is not too far from the one proposed by Srinivasan. It is a simple randomized algorithm that takes c as input and computes a suitable k for which it can assure that with high probability[2] the answer is going to be correct. Then, it proceeds with these three steps.

1. Color randomly all the nodes with $c/3$ colors.

2. Detect hyperedges whose constraints violate a (k, c)-coloring and re-color their nodes randomly with another set of $c/3$ colors.

3. Detect hyperedges whose constraints violate a (k, c)-coloring again, but now color them exhaustively with the last set of $c/3$ colors.

---

[2]With probability at least $1 - 1/n^{\alpha}$, for some constant α.

This algorithm can be used for obtaining any k, but in order to achieve the above mentioned guarantee, the values of c and k with which the steps are computed must satisfy a certain formula. Solving this formula, which involves also the maximum load $\mathcal{L}$ and the maximum degree Δ, allows the algorithm to compute a suitable k, which is chosen to be the minimum possible. Indeed, when the formula is satisfied, it is unlikely that we have many bad hyperedges after the second step, and thus the algorithm can be shown to run in polynomial time [Lu98]. As the k returned by the algorithm may be large for practical purposes, we exploit this idea in Section 6 to define a heuristic for the same problem.

5.2   Minimizing the number of wavelengths

In this case, for a given number of fibers k, the problem consists of finding the minimum number of colors c such that a valid (k, c)-coloring of the hypergraph exist. Noting that the number of colors is bounded by $\lceil n/k \rceil$ (and could be equal if the graph is a clique), we define a variable x_{ic} for each node and each color. It is given the value 1 when the node i is colored with color c and 0 otherwise. Then, the problem of minimizing c with k fixed can be solved by

Integer Program 2

$$
\begin{aligned}
\min \quad & \sum_c y_c && \textit{(minimize \# of colors)} \\
& \sum_c x_{ic} &=&\ 1 && \forall \textit{ node } i \\
& \sum_{i \in H} x_{ic} &\leq&\ k && \forall \textit{ color } c, \textit{hyperedge } H \\
& x_{ic} &\leq&\ y_c && \forall \textit{ color } c, \textit{node } i \\
& x_{ic}, y_c &\in&\ \{0, 1\} && \forall \textit{ color } c, \textit{node } i
\end{aligned}
$$

There are $O(n^2)$ variables and $O(n^2m)$ constraints (although it can be written with $O(nm)$ constraints if the solver generates cuts automatically).

The drawback of these IP formulations is that they are not symmetric and thus Branch-and-Bound will waste a lot of time iterating trough similar solutions [MT96]. The problem arises because after a variable is constrained by the algorithm, a permutation of them may still be feasible. Anyway, this problem can be solved using automatic pruning techniques, as described in [Mar01].

As the problems of minimizing k and minimizing c look at the same situation from different perspectives, an exact algorithm for one gives an exact algorithm that solves the other. For instance, if we only need a solution for a fixed k or a fixed c, using binary search we can give an answer to one problem using a routine for the other. In fact, the whole feasible region can be described by the *Pareto points* [Ste86], that is, the undominated pairs

of k and c that induce feasible colorings. For the case of approximations, that is not true anymore. Having an α-approximation algorithm for one problem does not imply the same for the other. This situation is evident for this pair of problems because one generalizes the coloring problem which is very hard and the other is a MIP that is not as hard to solve.

6 Implementation and Performance Evaluation

To computationally evaluate the problems, we implemented the two integer programs and the approximation algorithm, described in Section 5. This allowed us to evaluate the tradeoff between the performance and the running time of the exact version and the approximation. As we implemented an approximation algorithm for the problem of minimizing k, and computing exact solutions is equivalent for both problems, we used that problem to compare the results. We also report our findings in the experience of solving the problem of minimizing c.

The IPs for both versions of the problem were implemented using using *AMPL* [FGK93], while the approximation for the problem of minimizing k was implemented in C++. AMPL is a modeling language for mathematical programming that allows to type-in the IP formulations algebraically. In our case, the formulations that we used were exactly the ones described in the previous section. To solve them, AMPL uses a generic LP/IP solver, that in our case was *CPLEX*. Without making any effort in trying to tune up the parameters of the CPLEX Branch-and-Bound algorithm, the program could handle real-world instances in a relatively small PC (Pentium II, 500 MHz, 64MB).

The implementation of the approximate version was also straightforward. The three steps described in the Section 5 where implemented in a function that takes an hypergraph, a c and a k and tries to compute a (k,c)-coloring of the hypergraph. If a feasible coloring is found, it returns the number of fibers used and otherwise it fails. A difficulty we encountered is that the analysis guarantees a good behavior asymptotically, but our instances were too small. Indeed, the values of k that match the theoretical analysis of the algorithm are too large, even compared with the load, which is an upper bound. Hence, we sometimes obtained a set of nodes to be colored exhaustively in the third step which was larger than logarithmic in the size of the instance. In order to cope with this problem and obtain reasonable running time, we set a limit on this size. If the number of nodes to re-color was too large, the coloring procedure returned that the coloring was not feasible, which is likely to be the correct answer since there were only $\frac{c}{3}$ colors available.

The function just described was used as an oracle for the decision problem *"does a $(k,c)-coloring$ exist?"*. With it, we performed a binary search to minimize k for the given c. As the algorithm is randomized, we ran it 10 times per iteration and kept the best solution found so far.

We used instances based on an European network named *COST 239* [PCR00] and an American one [BCF^{+}00]. The European network joins 11 capitals using 24 multifiber links as depicted in Figure 2. As the referenced work included a demand matrix, we used it. Overall, there were 176 requests, which covered all possible pairs of cities. The maximum
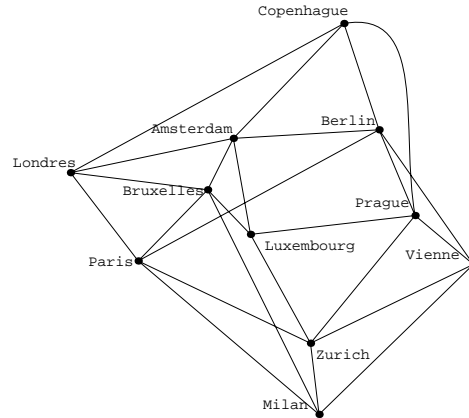
Figure 2: *European network COST 239.*
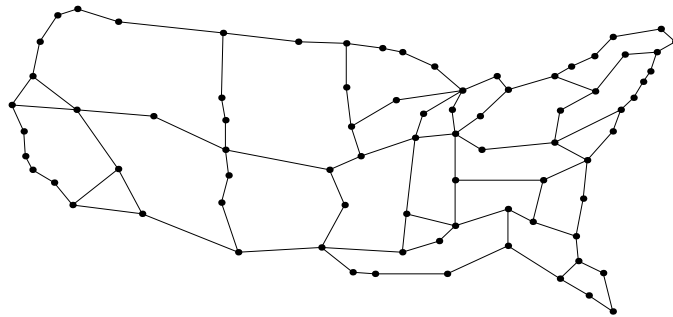
Figure 3: *American network.*

load of the routing was 29, which is also a lower bound for the number of colors in the single-fiber case.

The American network was bigger: it consists of 78 cities, interlinked by 102 arcs (see Figure 3). A demand matrix was not available in this case, and we had to generate it. For that, we used a gravitational model, for which we needed to associate weights to the cities. They represent, basically, the importance of every city. We made the weights proportional to the distance to 5 main population areas in the USA. Finally, the demand was generated between every two cities. The number of requests was made 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 report on a relatively big instance with 2022 requests and a load of 520.

Routing was implemented through a *minimum cost disjoint path problem* for each origin-destination pair. For each origin and destination, we computed the shortest total distance

c	$\min k$	time (sec.)	approx k	time (sec.)
1	29	< 1	29	< 1
3	10	< 1	29	< 1
9	4	< 1	9	< 1
15	2	2	6	< 1
33	1	< 1	3	< 1
63	1	< 1	2	< 1

Table 1: *Output for the European network*

of two disjoint paths linking them and distributed all the demand among the two paths. Routing the requests in such a way ensures that short paths are selected while maintaining two disjoint routes from each origin to each destination, which helps to improve the reliability.

When solving the exact version of the minimization of the number of fibers, the solver found feasible solutions reasonably fast. Except for the biggest instances (the American network with many colors), 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 $\mathcal{NP}$-hard. Nevertheless, this does not seem to be an issue for the instances generated from real-world networks.

6.1   Results

The number of fibers needed, as a function of the number of colors available, is displayed in Tables 1 and 2. We include instances with the number of colors increasing exponentially, rounded to the nearest multiple of 3. For the results of our heuristic, we report the best of 10 runs and the total computation time. Figure 4 shows the exact Pareto curve for the European instance. This kind of output could be of great utility when analyzing the tradeoffs involved between the necessary number of colors and fibers in reality.

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

The biggest dependency of the running time of the Integer Program 2, that 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 much time 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

c	$\min k$	time (sec.)	approx k	time (sec.)
1	520	< 1	520	< 1
3	174	6	520	554
9	58	228	173	479
15	35	2033	106	468
33	20	2.6 h.	47	419
63			29	368
129			15	299
255			7	272
513			6	276
1023			4	257

Table 2: *Output for the American network*
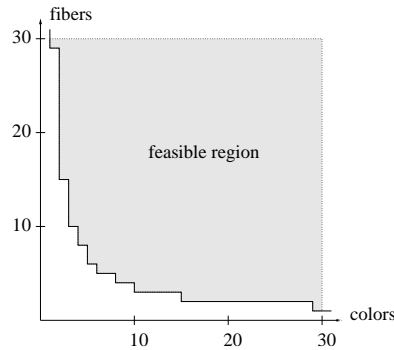
Figure 4: *Tradeoff between the number of colors and the number of fibers in the European network*

without altering the solution), the enumeration of the nodes of the Branch-and-Bound tree could not be completed in general. In any case, we had a proof of optimality. Indeed, 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.

7   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 since its straightforward IP formulation gave optimal solutions reasonably fast.

On the other hand, in all our real-world instances we found $w_k = \lceil \mathcal{L}/k \rceil$. Hence, the efficiency gain due to multifiber flexibility was not observed. However, since ad-hoc con-

structions prove that this gain can be enormous, practical instances could still be found where such a gain appear.

Furthermore, the heuristic that we implemented turned out to be very fast, but did not perform very well, despite the provably good asymptotic properties of the underlying randomized approximation algorithm. Therefore, we are currently working on the design of other heuristics for hypergraph coloring.

Another interesting research direction is to address 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.

References

[ACKP00] V. Auletta, I. Caragiannis, C. Kaklamanis, and P. Persiano. Randomized path coloring on binary trees. In Lecture Notes in Computer Science, editor, *3rd International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'00)*, volume 1913, pages 60–71. Springer-Verlag, 2000.

[BBG+97] 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, April 1997. IEEE Press.

[BBGK99] 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):743–754, May 1999.

[BCF+00] C. Blaizot, I. Cerutti, A. Fumagalli, M. Tacca, L. Valcarenghi, R. Jagannathan, A. Lardies, F. Masetti, and M. Garnot. Design and planning of transport networks based on SDH/SONET and WDM optimcal networks: methods, tools and applications. *Optical Networks Magazine*, pages 55–65, July 2000.

[Bea99] B. Beauquier. All-to-All Communication for some Wavelength-Routed All-Optical Networks. *Networks*, 33(3):179–187, 1999.

[BFM98] H. Bourdin, A. Ferreira, and K. Marcus. A performance comparison between graph and hypergraph topologies for passive star WDM lightwave networks. *Computer Networks and ISDN Systems*, 30:805–819, 1998.

[BGP+96] J-C. Bermond, L. Gargano, S. Pérennes, A. A. Rescigno, and U. Vaccaro. Efficient collective communication in optical networks. In *23rd International Colloquium on Automata, Languages and Programming (ICALP'96)*, volume 1099 of *Lecture Notes in Computer Science*, pages 574–585. Springer-Verlag, July 1996.

[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.

[CFK+01] I. Caragiannis, A. Ferreira, C. Kaklamanis, S. Pérennes, and H. Rivano. Fractional path coloring with applications to WDM networks. In *28th International Colloquium on Automata, Languages and Programming (ICALP'01)*, Lecture Notes in Computer Science, pages 732–743, Crete, Greece, July 2001. Springer-Verlag.

[CFM00a] D. Coudert, A. Ferreira, and X. Muñoz. A multihop-multi-ops optical interconnection network. *IEEE/OSA Journal of Lightwave Technology*, 18(12):2076–2085, 2000.

[CFM00b] D. Coudert, A. Ferreira, and X. Muñoz. Topologies for optical interconnection networks based on the optical transpose interconnection system. *OSA Applied Optics – IP*, 39(17):2965–2974, June 2000.

[CGK92]   I. Chlamtac, A. Ganz, and G. Karmi. Lightpath communications: An approach to high bandwidth optical WAN's. *IEEE Transactions on Communications*, 40(7):1171–1182, July 1992.

[EJK⁺99] T. Erlebach, K. Jansen, C. Kaklamanis, M. Mihail, and P. Persiano. Optimal wavelength routing on directed fiber trees. *Theoretical Computer Science*, 221(1-2):119–137, 1999.

[FGK93]   R. Fourer, D. M. Gay, and B. W. Kernighan. *AMPL: A Modeling Language for Mathematical Programming*. Duxbury Press / Brooks / Cole Publishing Company, 1993.

[Hoc97]   D. S. Hochbaum, editor. *Approximation Algorithms for NP-Hard Problems*. PWS Publishing Compagny, 1997.

[HV99]   E. Hyytiä and J. Virtamo. Wavelength assignment in multifibre in WDM-networks. Technical Report COST257TD(99)04, Helsinki University of Technology, February 1999.

[KK01]   J. Kleinberg and A. Kumar. Wavelength Conversion in Optical Networks. *Journal of algorithms*, 38(1):25–50, January 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.

[LS99]   L. Li and A. Somani. Blocking performance analysis of fixed-paths least-congestion routing in multifiber wdm networks. In *All optical networking 1999: architecture, control, and management issues*, pages 56–67, Boston, MA, 1999.

[LS00]   G. Li and R. Simha. On the Wavelength Assignement Problem in Multifiber WDM Star and Ring Networks. *IEEE Infocom*, 3:1771–1780, 2000.

[Lu98]   C. Lu. Deterministic hypergraph coloring and its applications. In *Proceedings of the 2nd International Workshop on Randomization and Approximation Techniques in Computer Science*, volume 1518 of *Lecture Notes in Computer Science*, pages 35–46. Springer-Verlag, October 1998.

[Mar01]   F. Margot. Pruning by isomorphism in Branch-and-Cut. In K. Aardal and B. Gerards, editors, *Lecture Notes in Computer Science 2081: 8th IPCO Conference*, pages 304–317. Springer, June 2001.

[MS00]   L. Margara and J. Simon. Wavelength assignment problem on all-optical net-
           works with k fibres per link. In *27rd International Colloquium on Automata,
           Languages and Programming*, pages 768–779, 2000.

[MS01]   L. Margara and J. Simon. Decidable properties of graphs of all-optical networks.
           In *28rd International Colloquium on Automata, Languages and Programming*,
           pages 768–779, 2001.

[MT96]   A. Mehrotra and M. A. Trick. A column generation approach for graph coloring.
           INFORMS Journal on Computing, 8(4):344–354, 1996.

[PCR00]   INRIA MASCOTTE Project, Alcatel Corporate Research Centre, and
           France Télécom R&D. Délivrable 1. Document soumis à autorisation légale
           de France Télécom R&D, February 2000. Projet RNRT PORTO, Planification
           et Optimisation des Réseaux de Transport Optiques.

[RM88]   B. Ramamurty and B. Mukherjee. Wavelength conversion in WDM networking.
           IEEE Journal on Selected Areas in Communications, 16(7):1061–1073, Septem-
           ber 1988.

[RS95]   R. Ramaswami and K. N. Sivarajan. Routing and wavelength assignment in all-
           optical networks. *IEEE/ACM Transactions on Networking*, 3(5):489–500, 1995.

[RS98]   R. Ramaswami and K. N. Sivarajan. *Optical Networks: a practical perspective.*
           Morgan-Kaufman, 1998.

[Sri96]   A. Srinivasan. An extension of the Lovász Local Lemma, and its applications to
           integer programming. In *ACM-SIAM Symposium on Discrete Algorithms*, pages
           6–15, 1996.

[Ste86]   R. E. Steuer. *Multiple Criteria Optimization: Theory, Computation and Appli-
           cation.* Wiley, New York, 1986.

[ZQ98]   X. Zhang and C. Qiao. Wavelength assignment for dynamic traffic in multi-fiber
           WDM networks. In *ICCCN'98*, pages 479–585, 1998.

Contents