

Limited Bandwidth in Multiple-Fiber All-Optical Caterpillars: a Minimization Problem

Thomas Erlebach¹, Aris Pagourtzis², Katerina Potika², and Stamatis Stefanakos¹

¹ Computer Engineering and Networks Laboratory (TIK), ETH Zürich, 8092 Zürich, Switzerland, Email: {erlebach,stefanak}@tik.ee.ethz.ch ***

² Department of Computer Science, School of ECE, National Technical University of Athens, 15780 Zografou, Greece, Email: {pagour,epotik}@cs.ntua.gr †

Abstract. In multiple-fiber WDM networks it is important to minimize the number of active fiber links that need to be allocated in order to satisfy a given communication demand. Minimizing this number results in reduced network cost. Equivalently, one wants to minimize the number of wavelength collisions on each link of the network. This is modeled as the MINIMUM COLLISIONS PATH MULTICOLORING problem, which differs from the classical path coloring problems in that using the same colors for overlapping paths is now allowed. More formally, we are given a graph $G = (V, E)$, a set of communication requests (pairs of nodes) and a number of available colors. Requests can be undirected or directed, corresponding to full-duplex or one-way communication respectively. Each request must be assigned a path and a color, and the multiplicity of fibers per edge is determined by the maximum number of paths colored with the same color among paths that use this edge. The goal is to minimize the sum of fiber multiplicities over the network. The caterpillar topology is of particular interest because it often describes real networks or parts of a real network. In this paper we present an efficient approximation algorithm for caterpillar networks with undirected requests; the algorithm achieves approximation ratio $1 + \frac{5|E^*|}{OPT}$ (≤ 6 in the worst case), where E^* is the set of edges used by at least one communication path. By appropriate adaptation of our algorithm we obtain an approximation ratio $1 + \frac{4|E^*|}{OPT}$ for the case of directed requests. These ratios get closer to 1 as the network traffic increases.

1 Introduction

Optical networks are playing a vital role in today's communication, such as the Internet and other data services, since they offer high bandwidth, low delay and low bit-error rates. A network is called all-optical, if every signal remains in optical form

*** T. Erlebach and S. Stefanakos were partially supported by the Swiss National Science Foundation (project AAPCN) and by EU Thematic Network APPOL II, IST-2001-32007, with funding by the Swiss Federal Office for Education and Science.

† A. Pagourtzis and K. Potika were partially supported by the Human Potential Programme of EU, contract no HPRN-CT-1999-00104 (AMORE), while visiting the Institute of Theoretical Computer Science, ETH Zürich (A. Pagourtzis) and the University of Konstanz (K. Potika).

from its transmitter to its receiver without changing to electrical form. The technology that enables transmitting more than a single wavelength down the same optical fiber is called *Wavelength Division Multiplexing (WDM)*; it allows many signals to be transferred through the same fiber, each at the rate of several Gigabits per second, under the constraint that each of them uses a different wavelength from a set of w fixed wavelengths.

A connection request is a pair (transmitter, receiver) and is satisfied, when a path from its transmitter to its receiver is specified and a wavelength is assigned to all links of that path. Placing intermediate signal conversion from/to electrical form lifts the restriction of using the same wavelength on every link of the path, but such devices decrease the performance of the network and therefore are not dealt with in our approach. On the other hand, all-optical wavelength converters do not have this disadvantage but are very expensive, hence their use is undesirable.

New technologies make it possible to install more fibers in each link. This has given rise to the study of new problems related to routing and wavelength assignment in WDM all-optical networks. In this approach, signals that pass through the same link may use the same wavelength provided that they traverse different parallel fibers. Such multi-fiber all-optical networks provide switching capabilities with the help of optical crossconnects, where signals that are on the same fiber of an incoming link may continue on different fibers of an outgoing link without changing wavelength, and vice versa.

The optimization problem that we study is the minimization of the number of active fibers that need to be allocated in order to satisfy all given communication requests. Reserving only the number of necessary (active) fibers on each link is much more cost effective than reserving for the whole network the same number of parallel fibers on all links. Naturally, the number of active fibers on a link of the network is the Maximum number of signals that use the same wavelength among signals that pass through this link.

The problem is usually referred to in the literature as the MINIMUM COLLISIONS PATH MULTICOLORING (MINCOLLISIONS-PMC) problem [16, 14, 7]: We are given a graph $G = (V, E)$, a set of communication requests (pairs of nodes) and a number of available colors. Each request must be assigned a path and a color, and the multiplicity of fibers per edge is determined by the maximum number of color collisions (repetitions) among paths that use this edge. The goal is to minimize the sum of fiber multiplicities over the network. MINCOLLISIONS-PMC differs from the classical path coloring problem in that we now allow overlapping paths to share a color.

A *caterpillar* is a network that consists of a path (chain) called the *backbone*, and several other chains, which may intersect only on the backbone. The caterpillar topology is of particular interest because it often describes real networks or parts of a real network. MINCOLLISIONS-PMC is clearly NP-hard for caterpillar networks since star networks are a special case of caterpillars and the problem is already known to be NP-hard for stars [16].

1.1 Our Contribution

In this work, we first present an algorithm for MINCOLLISIONS-PMC in caterpillars (undirected requests) for the special case of paths that use at least one *backbone edge*; the algorithm uses a transformation to bipartite edge coloring and finds solutions with a total number of fibers that is at most $OPT + 3|E^*|$, where OPT denotes the number of active fibers of an optimal solution and E^* is the set of edges that are used by at least one path. Next, we extend the algorithm to any communication pattern using at most $OPT + 5|E^*|$ fibers in total. For caterpillars with directed requests we present an algorithm that uses at most $OPT + 4|E^*|$ fibers totally in both directions, where E^* now denotes the set of *directed* edges that are used by at least one path.

Clearly, $OPT \geq |E^*|$, hence the approximation ratios achieved by the above algorithms are bounded by 4, 6 and 5 respectively; although these ratios are not better in the worst case than the ratio 4, achieved by an algorithm for MINCOLLISIONS-PMC in trees which is implicit in [2] (see [7]), they prove to be much better if the network traffic is heavy. In particular, they tend to 1 as the load of the network increases.

1.2 Related Work

The MINCOLLISIONS-PMC problem was first studied in [16]. It was shown that it can be solved optimally in polynomial time in chain networks, while for star and ring networks 2-approximation algorithms were given. A second algorithm for MINCOLLISIONS-PMC in chain networks was presented in [21].

For general tree networks, Erlebach et al. have given two Approximation algorithms with ratios $1 + 4|E| \log |V| / OPT$ and 4 respectively [7]. Comparing those results with the results of this paper one can see that the algorithms presented here guarantee no better approximation ratios in the worst case. However, for heavily loaded instances the approximation ratios of the new algorithms are much better.

Nomikos et al. [14, 15] study a generalization of MINCOLLISIONS-PMC with different link costs and present constant-ratio approximation algorithms for rings and spiders.

The problem where the number of parallel fibers is given and the goal is to satisfy all requests using a minimum number of wavelengths, is called MINIMUM COLORS PATH MULTICOLORING (MINCOLORS-PMC) [7]. If, in addition, the number of parallel fibers μ is the same for all edges the problem is called MINIMUM COLORS UNIFORM PATH MULTICOLORING (MINCOLORS-UPMC). For MINCOLORS-PMC in trees, a 4-approximation algorithm both for the undirected and the directed case was given recently by Chekuri, Mydlarz and Shepherd [2]. MINCOLORS-UPMC is studied in [11, 12] for stars and rings, where approximation algorithms with small constant ratios are given; in [8], a new approach for general network topologies is developed, based on conflict hypergraphs. Li and Simha [10] have also studied MINCOLORS-UPMC for undirected trees and observed that a valid multicoloring using at most $\lceil 3L / (2\mu) \rceil$ colors can be computed efficiently using an algorithm for the single-fiber problem (PC) due to Raghavan and Upfal [19]. The currently best known algorithm for PC in undirected trees achieves an asymptotic approximation ratio of 1.1. This follows from the algorithm of Nishizeki and Kashiwagi [13] for edge coloring of multigraphs since approximation

algorithms for PC in undirected trees and edge coloring of multigraphs are interchangeable [9, 5]. In directed trees, the best known algorithm is a $5/3$ -approximation (using at most $5L/3$ colors) due to Erlebach et al. [6].

Another interesting related problem is the MAXIMUM ROUTING AND PATH COLORING PROBLEM (MAXRPC). An instance of MAXRPC is similar to an instance of MINCOLLISIONS-PMC, but no color repetitions on paths that use the same edge are allowed; the goal is to route and color as many requests as possible with the available colors. The MAXRPC problem has been studied for several topologies. In chains, also known as the “ k -coloring of intervals” problem, it can be solved exactly [1]. Nomikos, Pagourtzis and Zachos [17] have considered the problem in ring networks and have proposed a $3/2$ -approximation algorithm for the undirected case and a $11/7$ -approximation algorithm for the directed case; they have also given a $3/2$ -approximation algorithm for rings with (undirected or directed) pre-routed requests [18]. For trees, a 1.58 approximation for the undirected case was presented by Wan and Liu [20], while for the directed case a 2.22-approximation is due to Erlebach and Jansen [4].

2 Notation - Preliminaries

Let \mathcal{P} be a set of paths on a graph G ; the load $L(e, \mathcal{P})$ of edge e w.r.t. \mathcal{P} is the number of paths in \mathcal{P} that go over e . The load $L(\mathcal{P}) = \max_{e \in E} L(e, \mathcal{P})$ of G is the maximum number of paths that go over the same edge of G . We will omit \mathcal{P} and simply write $L(e)$ and L whenever it is clear from the context which set \mathcal{P} we mean. A coloring of a set of paths where color collisions are allowed among paths that share an edge is called *path multicoloring*, as opposed to *path coloring* which usually means that paths with common edges must receive different colors. For a set of connection requests on a multi-fiber network, a multicoloring of the corresponding paths determines a wavelength assignment on multiple fibers in the following sense: requests routed over edge e whose paths receive the same color must be assigned the same wavelength and therefore must be routed through different fiber links; therefore, the *maximum number of color collisions* on e , denoted by $\mu(e)$, determines the number of parallel fiber links needed on edge e to establish the corresponding connections. If we have a limited number w of available wavelengths per fiber a reasonable task is to try to reduce the number of fiber links used on the network. The corresponding minimization problem is:

MINIMUM COLLISIONS PATH MULTICOLORING (MINCOLLISIONS-PMC). *Given a graph $G = (V, E)$, a set of paths \mathcal{P} on G and a number w , find a path multicoloring of \mathcal{P} with w colors such that $\sum_e \mu(e)$ is minimized, where $\mu(e)$ denotes the maximum number of times that any color appears on edge e .*

An algorithm A for a minimization problem Π is a ρ -approximation if for every instance I of Π , A runs in time polynomial in $|I|$ and delivers a solution with objective value $SOL(I)$ at most $\rho \cdot OPT(I)$, where $OPT(I)$ denotes the objective value of an optimal solution for I . If the problem and the instance under consideration are clear from the context, we will simply use OPT to denote the objective value of an optimal solution for the instance and SOL to denote the objective value of the solution returned

by A . For our problem (MINCOLLISIONS-PMC) we will sometimes use $OPT(\mathcal{P})$ and $SOL(\mathcal{P})$ to denote the corresponding quantities for a specific set of paths \mathcal{P} .

Note that a natural lower bound for the maximum number of color collisions on an edge e is $\mu(e) \geq \lceil L(e, \mathcal{P})/w \rceil$. Hence, for any instance (G, \mathcal{P}, w) of MINCOLLISIONS-PMC it holds: $OPT(\mathcal{P}) \geq \sum_{e \in E} \lceil L(e, \mathcal{P})/w \rceil$.

A *chain* (or path) is a connected graph with internal nodes of degree 2 and two terminal nodes of degree 1. A *star* is a tree graph with only one node of degree > 1 . A *spider*, also known as generalized star, is a tree graph with only one node of degree > 2 (called the *center*); it can also be seen as a star whose edges have been replaced by chains.

A *caterpillar* is a tree graph that consists of several spiders whose centers lie on a chain (called *backbone*, or spine). Equivalently, a caterpillar consists of a primary chain (the backbone) together with several secondary chains (the *legs*) attached to the backbone by coalescing a terminal of each leg with a backbone node. An example of a caterpillar is given in Figure 1.

Let $G = (V, E)$ be a caterpillar graph; the set of nodes V is partitioned into two sets of nodes, $V_b = \{v_0, v_1, \dots, v_{n'-1}\}$ (set of backbone nodes) and $V_l = \{u_1, u_2, \dots, u_{n-n'}\}$ (set of leg nodes), and the set of edges E is partitioned into two sets of edges, $E_b = \{\{v_0, v_1\}, \{v_1, v_2\}, \dots, \{v_{n'-2}, v_{n'-1}\}\}$ and $E_l = E \setminus E_b$. Note that the backbone of G is the chain $G_b = (V_b, E_b)$; the graph obtained by removing the backbone edges E_b , i.e. $G_l = (V, E_l)$, is a forest of spiders.

For the purposes of our problem, we are only interested in a subset $E^* \subseteq E$ of edges that are used by at least one path, and the corresponding subsets of backbone and leg edges, E_b^* and E_l^* respectively; for convenience, we will often omit the “*” symbol.




Fig. 1. Example of a caterpillar.

Nodes in V_l are indexed so that nodes that belong to the same spider take consecutive indices. The indices of the spider of v_i are smaller than the indices of the spider of v_j for $i < j$. The same holds for each leg separately starting from the tip (outermost node) of the leg ? see Figure 1. We also use the following convention: the index of any node of a spider is considered smaller than the index of the center of the spider v_i

and larger than the index of v_{i-1} (the center of the “previous” spider). In the case of undirected requests and paths we say that a path *starts* at its smaller-index endpoint and *finishes* at its larger-index endpoint.

We consider a partition of the path set \mathcal{P} into a set \mathcal{P}_b that contains the paths that go over at least one backbone edge, from now on *backbone paths*, and a set \mathcal{P}_l that contains the remaining paths ($\mathcal{P}_l = \mathcal{P} \setminus \mathcal{P}_b$). Note that each path in \mathcal{P}_l lies exclusively on a single spider of G_i ; hence, we call paths in \mathcal{P}_l *spider paths*.

For a set of paths $\mathcal{Q} \subseteq \mathcal{P}_b$, we denote by $L_s(v_i, \mathcal{Q})$ the number of paths in \mathcal{Q} that start at node v_i (or at any node in the spider of v_i) and by $L_f(v_i, \mathcal{Q})$ the number of paths in \mathcal{Q} that end at node v_i (or at any node in the spider of v_i).

3 Undirected (Full-Duplex) Communication

In full-duplex communication it is assumed that a two-way connection is established between the given endpoints, and signals in both directions use the same set of edges and use the same wavelength. One can use undirected paths to represent such connections. In this section we deal with this case in caterpillar networks.

3.1 An Algorithm for Multi-Coloring Backbone Paths

We will first present an algorithm for MINCOLLISIONS-PMC in caterpillars with only backbone paths. The algorithm (Procedure 1) works as follows: we transform an instance of our problem to an instance of edge coloring of a bipartite multigraph. Coloring the edges of the constructed bipartite multigraph yields a multicoloring for the paths of our initial instance.

In steps 1 and 2 of Procedure 1 we put each path into one “starting” and one “finishing” group of paths. Dummy paths are added in order to make the groups *full*; a full group is one that contains exactly w paths (see Figure 2).

We define the following operation for $a, b \in \mathbf{N}, b > 0$

$$a \overline{\text{mod}} b = (-a) \text{ mod } b = \begin{cases} b - a \text{ mod } b & \text{if } (a \text{ mod } b) \neq 0 \\ 0 & \text{if } (a \text{ mod } b) = 0 \end{cases}$$

In other words, $a \overline{\text{mod}} b$ is the minimum number that needs to be added to a in order to obtain an exact multiple of b .

We will prove that Procedure 1 finds a multicoloring for paths in \mathcal{P}_b such that the number of color repetitions of any color on an edge e is at most $\lceil L(e, \mathcal{P}_b)/w \rceil + 3$, i.e. at most three away from the optimal. We first show the following lemma:

Lemma 1. *After step 2 of Procedure 1, for any backbone node $v_i, 0 \leq i < n'$, it holds that the set of paths that start (finish) at the spider centered at v_i is partitioned into full starting (resp. finishing) groups.*

Proof. The lemma clearly holds for the original backbone paths (i.e. in \mathcal{P}_b), together with a small ($< w$) number of dummy paths, due to the construction of the groups.




Fig. 2. Filling, grouping and joining paths that start or finish at a node v_i . It is assumed that the number of available colors is 3. Dummy paths are shown by dotted lines. The small rectangle depicts a join.

Procedure 1 *Multi-coloring the Backbone Paths*

Input: caterpillar network G , set of paths \mathcal{P}_b , number of colors w .

Output: path multicoloring of \mathcal{P}_b with w colors.

1. For each backbone edge $e \in E_b$, add to \mathcal{P}_b a number of $d(e) = w + L(e, \overline{\mathcal{P}_b}) \bmod w$ dummy (unit length) paths, consisting only of edge e . Let the new path set be called \mathcal{P}'_b .
 2. For each backbone node $v \in V_b$ do
 - (a) Let $\mathcal{P}_v \subseteq \mathcal{P}'_b$ be the subset of \mathcal{P}'_b that contains paths starting in the spider centered at v . Order paths in \mathcal{P}_v in ascending order of the index of their endpoint in the spider; paths starting at the center are placed after paths starting at the legs and dummy paths are the last among them. Partition \mathcal{P}_v into several groups of w paths and one group of at most w paths, following the order. Add each full group (i.e. consisting of w paths) to the set of starting groups S .
 - (b) Repeat step 2a for the subset of \mathcal{P}'_b (denoted by \mathcal{P}'_v) that contains paths finishing in the spider centered at v . Add the formed full groups to the set of finishing groups F .
 - (c) Join remaining dummy paths that finish at v with remaining dummy paths that start at v . Let the new path set be called \mathcal{P}''_b .
 3. Construct a bipartite graph $H = (S \cup F, A)$ as follows: The sets of nodes S, F are the set of starting groups and the set of Finishing groups respectively. For each path in \mathcal{P}''_b there is one edge in A connecting its starting and its finishing groups.
 4. Compute an edge coloring of H using an appropriate algorithm (e.g. the algorithm of Cole, Ost and Schirra [3], which will take $O(|\mathcal{P}''_b| \log w)$ time). Assign to every path in \mathcal{P}''_b the color of the corresponding edge in H .
-

Note that there are always enough paths so as to form full groups for all the original paths since at least w dummy paths are added on each edge.

It suffices to prove that the remaining finishing and starting (dummy) paths at a node v_i are equal in number and, therefore, after joining them there are no remaining paths. The set of paths \mathcal{P}'_b resulting after step 1 of Procedure 1 has the property that, on each edge, the load is an exact multiple of w ; hence, there are $k, l \in \mathbf{N}$ such that $L(e_i, \mathcal{P}'_b) = kw$ and $L(e_{i+1}, \mathcal{P}'_b) = lw$, where $e_i = \{v_{i-1}, v_i\}, 0 < i < n'$. Let also x be the number of paths that span both edges e_i and e_{i+1} . It holds:

$$L_f(v_i, \mathcal{P}'_b) = L(e_i, \mathcal{P}'_b) - x = kw - x \quad L_s(v_i, \mathcal{P}'_b) = L(e_{i+1}, \mathcal{P}'_b) - x = lw - x$$

Let us denote by $L'_f(v_i, \mathcal{P}'_b)$ (resp. $L'_s(v_i, \mathcal{P}'_b)$) the number of non-grouped finishing (resp. starting) paths at v_i , after constructing the full groups. It holds:

$$\begin{aligned} L'_f(v_i, \mathcal{P}'_b) &= L_f(v_i, \mathcal{P}'_b) \bmod w = (kw - x) \bmod w = \overline{x \bmod w} \\ &= (lw - x) \bmod w = L_s(v_i, \mathcal{P}'_b) \bmod w = L'_s(v_i, \mathcal{P}'_b) \end{aligned}$$

Hence, all non-grouped paths at v_i are joined into larger paths. \square

The following corollary is immediate, since all groups are full and there do not remain non-grouped paths.

Corollary 1. *After step 2 of Procedure 1, every path in \mathcal{P}''_b belongs to exactly one full starting group and exactly one full finishing group.*

The next two propositions state that the multicoloring obtained by Procedure 1 is not far from the lower bound for both the backbone and the leg edges.

Proposition 1. *Consider any edge $e_i \in E_b, 0 \leq i \leq n' - 1$. After executing Procedure 1 each color is repeated exactly $L(e_i, \mathcal{P}'_b)/w$ ($= \lceil L(e_i, \mathcal{P}_b)/w \rceil + 1$) times on paths that go over e_i .*

Proof. We first notice that the bipartite graph H has degree w (in fact it is w -regular) and therefore it can be edge-colored with w colors. Paths in a group correspond to adjacent edges in H , so they receive different colors, i.e. each color is used exactly once in each group.

We can now prove the lemma by induction on the index of edges.

The claim holds for e_1 since all paths on e_1 start at the spider centered in v_0 , therefore, by Lemma 1 they are partitioned into $L(e_1, \mathcal{P}''_b)/w$ groups. Hence, each color is used exactly $L(e_1, \mathcal{P}''_b)/w = L(e_1, \mathcal{P}'_b)/w$ times.

Now, if the claim holds for edge e_i it also holds for e_{i+1} ; to see why, let f_i (s_i) be the number of finishing (resp. starting) groups at v_i . Then, each color is used $r_i = L(e_i, \mathcal{P}'_b)/w + s_i - f_i$ times on e_{i+1} . It suffices to observe that:

$$r_i = \frac{L(e_i, \mathcal{P}'_b) + s_i w - f_i w}{w} = \frac{L(e_{i+1}, \mathcal{P}'_b)}{w}$$

The fact $L(e_i, \mathcal{P}'_b)/w = \lceil L(e_i, \mathcal{P}_b)/w \rceil + 1$ is immediate From the number of dummy paths added in step 1 of Procedure 1. \square

Proposition 2. *Consider any edge $e \in E_l$. After executing Procedure 1 each color is repeated at most $\lceil L(e, \mathcal{P}_b)/w \rceil + 3$ times on paths that go over e .*

Proof. Let \mathcal{S} be the spider where e belongs. The load of paths $\in \mathcal{P}_b$ that go over e can be divided in two quantities: the load $L_s(e, \mathcal{P}_b)$ of paths that start at \mathcal{S} and the load $L_f(e, \mathcal{P}_b)$ of paths that finish at \mathcal{S} . Due to the way paths are grouped in step 2 paths that start (finish) at \mathcal{S} and go over e are grouped consecutively, therefore each color is used at most $\lceil L_s(e, \mathcal{P}_b)/w + 1 \rceil$ ($\lceil L_f(e, \mathcal{P}_b)/w + 1 \rceil$, resp.) times among these paths. (The “+1” term is due to the fact that some paths may be used to fill a group that starts in some previous edge and some other paths may be used to start a group that continues in some following edge.)

Altogether, the number of color repetitions on e is:

$$\begin{aligned} SOL(e, \mathcal{P}_b) &\leq \lceil L_s(e, \mathcal{P}_b)/w \rceil + 1 + \lceil L_f(e, \mathcal{P}_b)/w \rceil + 1 \leq \\ &\quad \left\lceil \frac{L_s(e, \mathcal{P}_b) + L_f(e, \mathcal{P}_b)}{w} \right\rceil + 3 = \left\lceil \frac{L(e, \mathcal{P}_b)}{w} \right\rceil + 3 \end{aligned}$$

□

Theorem 1. *Procedure 1 is a $(1 + 3|E|/OPT)$ -approximation algorithm For the problem MINCOLLISIONS-PMC in caterpillars with only backbone paths.*

Proof. Recall that for each edge e and set of paths \mathcal{P}_b the minimum number of fibers needed is $OPT(e, \mathcal{P}_b) \geq \lceil L(e, \mathcal{P}_b)/w \rceil$.

From Propositions 1 and 2 we get bounds for color multiplicity on backbone and leg edges. Adding over the whole caterpillar we have:

$$\begin{aligned} SOL(\mathcal{P}_b) &= \sum_{e \in E_b} SOL(e, \mathcal{P}_b) + \sum_{e \in E_l} SOL(e, \mathcal{P}_b) \\ &\leq \sum_{e \in E_b} (OPT(e, \mathcal{P}_b) + 1) + \sum_{e \in E_l} (OPT(e, \mathcal{P}_b) + 3) \\ &\leq OPT(\mathcal{P}_b) + |E_b| + 3|E_l| < OPT(\mathcal{P}_b) + 3|E| \end{aligned}$$

Hence, $SOL(\mathcal{P}_b) < (1 + \frac{3|E|}{OPT(\mathcal{P}_b)}) \cdot OPT(\mathcal{P}_b)$.

□

3.2 Approximating MinCollisions-PMC in Caterpillars

The basic idea is to combine the algorithm that multi-colors paths in \mathcal{P}_b (backbone paths) with an algorithm due to Nomikos et al. [14] that can be used to multi-color paths on a spider and hence paths in \mathcal{P}_l (spider paths). It holds:

Lemma 2 ([14]). *There is an algorithm for MINCOLLISIONS-PMC in spiders which, given a set of undirected paths \mathcal{Q} on a spider \mathcal{S} , produces a multicoloring with no more than $\lceil L(e, \mathcal{Q})/w \rceil + 1$ color collisions on any edge e of \mathcal{S} .*

The algorithm for caterpillars follows.

Algorithm 1 *Minimum Collisions Path Multi-coloring in Caterpillars*

1. Partition \mathcal{P} into two subsets \mathcal{P}_b (backbone paths) and \mathcal{P}_l (spider paths).
 2. Multi-color paths $\in \mathcal{P}_b$ using Procedure 1.
 3. Multi-color paths $\in \mathcal{P}_l$ using the algorithm for MINCOLLISIONS-PMC in spiders from [14].
-

Theorem 2. *Algorithm 1 is a $(1+5|E|/OPT)$ -approximation algorithm for MINCOLLISIONS-PMC in caterpillars.*

Proof. We first compute upper bounds for the number of color collisions on backbone edges and on leg edges.

- For any edge e in E_b the load is only due to paths in \mathcal{P}_b . Hence, Proposition 1 applies and $SOL(e, \mathcal{P}) = SOL(e, \mathcal{P}_b) \leq OPT(e, \mathcal{P}_b) + 1 = OPT(e, \mathcal{P}) + 1$.
- For any edge e in E_l the load can be analyzed as the load from paths in \mathcal{P}_b plus the load from paths in \mathcal{P}_l . Combining Proposition 2 with Lemma 2 we get:

$$\begin{aligned}
SOL(e, \mathcal{P}) &\leq SOL(e, \mathcal{P}_b) + SOL(e, \mathcal{P}_l) \\
&\leq \left\lceil \frac{L(e, \mathcal{P}_b)}{w} \right\rceil + 3 + \left\lceil \frac{L(e, \mathcal{P}_l)}{w} \right\rceil + 1 \\
&\leq \left\lceil \frac{L(e, \mathcal{P}_b) + L(e, \mathcal{P}_l)}{w} \right\rceil + 5 \\
&= \left\lceil \frac{L(e, \mathcal{P})}{w} \right\rceil + 5 \leq OPT(e, \mathcal{P}) + 5
\end{aligned}$$

Adding over the whole caterpillar we have:

$$\begin{aligned}
SOL(\mathcal{P}) &= \sum_{e \in E_b} SOL(e, \mathcal{P}) + \sum_{e \in E_l} SOL(e, \mathcal{P}) \\
&\leq \sum_{e \in E_b} (OPT(e, \mathcal{P}) + 1) + \sum_{e \in E_l} (OPT(e, \mathcal{P}) + 5) \\
&\leq OPT(\mathcal{P}) + |E_b| + 5|E_l| < OPT(\mathcal{P}) + 5|E|
\end{aligned}$$

Hence, $SOL(\mathcal{P}) < (1 + \frac{5|E|}{OPT(\mathcal{P})}) \cdot OPT(\mathcal{P})$. □

4 Directed (One-Way) Communication

The case of one-way communication can be modeled by directed requests and paths. On every edge e each direction is considered separately meaning that the total number of fibers needed for e is the sum of the maximum number of color repetitions in each direction of e .

We partition again the directed path set \mathcal{P} into two subsets \mathcal{P}_b (backbone paths) and \mathcal{P}_l (spider paths). We denote by $L_{in}(e)$ the load on the one direction of edge e (i.e.

from its smaller-index endpoint to its larger-index endpoint¹) and by $L_{out}(e)$ the load on the other direction of e .

Let E_{1b} (E_{1l}) be the subset of edges E_b (E_l respectively) that have paths only in the one of the two directions and E_{2b} (E_{2l} respectively) be the set of edges with paths in both directions. Also let $E_2 = E_{2b} \cup E_{2l}$ and $E_1 = E_{1b} \cup E_{1l}$.

Note that, if we consider the underlying graph as being bidirected (i.e. there are two opposite directed edges instead of each undirected one) and denote by E^* the set of directed edges that are used by at least one path it holds: $|E^*| = 2|E_2| + |E_1|$. Since any algorithm will use at least one fiber for each directed edge it holds: $OPT \geq |E^*|$.

4.1 Multi-Coloring Directed Backbone Paths

In the directed case directions are given, thus in every edge paths using both directions may exist. We multicolor the backbone paths in \mathcal{P}_b using independently Procedure 1 for each direction of the backbone. This is possible because each path p in \mathcal{P}_b has the same direction on all its backbone edges.

From Proposition 1 we get the following corollary:

Corollary 2. *Consider any directed edge $e \in E_b$. After calling Procedure 1 once for each direction, each color is repeated at most $\left\lceil \frac{L_i(e, \mathcal{P}_b)}{w} \right\rceil + 1$ times on paths that go over e in each direction $i \in \{in, out\}$.*

For the edges that belong to the legs of the spiders we have the following:

Proposition 3. *Consider any edge $e \in E_l$. After executing Procedure 1 once for each direction of the backbone, each color is repeated at most $\left\lceil \frac{L_i(e, \mathcal{P}_b)}{w} \right\rceil + 3$ in each direction $i \in \{in, out\}$ of edge e .*

Proof. Paths in \mathcal{P}_b using $e \in E_l$ in the same direction may have different directions in the backbone edges and that way the paths on edge $e \in E_l$ are partitioned into two sets therefore solved independently using Procedure 1. For each subset it can be shown, using a similar argument as that on the proof of Proposition 2, that each color is repeated at most $\left\lceil \frac{L_x}{w} \right\rceil + 1$, where L_x is the load of this subset on e .

Let $L'_i(e, \mathcal{P}_b)$ ($L''_i(e, \mathcal{P}_b)$) denote the number of paths that have direction i on e and “in” on backbone edges (respectively “out”). Therefore for each direction $i \in \{in, out\}$:

$$SOL_i(e, \mathcal{P}_b) \leq \left\lceil \frac{L'_i(e, \mathcal{P}_b)}{w} \right\rceil + 1 + \left\lceil \frac{L''_i(e, \mathcal{P}_b)}{w} \right\rceil + 1 \leq \left\lceil \frac{L_i(e, \mathcal{P}_b)}{w} \right\rceil + 3$$

□

¹ Recall that, by convention, nodes on each leg of a spider are considered smaller-indexed than the center of the spider.

4.2 MinCollisions-PMC in Caterpillars with Directed Paths

We combine the algorithms that multicolor the backbone paths (\mathcal{P}_b) and the spider paths (\mathcal{P}_l). Note that the algorithm presented in [14], originally designed for undirected paths in spiders, yields optimal solutions if applied to spiders with directed paths.

Lemma 3 ([14]). *A multicoloring of a directed path set \mathcal{Q} in a spider such that no color is repeated more than $\left\lceil \frac{L_i(e, \mathcal{Q})}{w} \right\rceil$ times in direction $i \in \{\text{in}, \text{out}\}$ of edge e can be computed in polynomial time.*

Theorem 3. *The algorithm for MINCOLLISIONS-PMC in caterpillars with directed paths is a $1 + \frac{4 \cdot |E^*|}{OPT}$ -approximation algorithm (a 5-approximation algorithm in the worst case).*

Proof. For each edge $e \in E_l$ in each direction $i \in \{\text{in}, \text{out}\}$ we have (by Proposition 3 and Lemma 3):

$$SOL_i(e, \mathcal{P}) \leq \left\lceil \frac{L_i(e, \mathcal{P}_b)}{w} \right\rceil + 3 + \left\lceil \frac{L_i(e, \mathcal{P}_l)}{w} \right\rceil \leq \left\lceil \frac{L_i(e, \mathcal{P})}{w} \right\rceil + 4 \quad (1)$$

Hence, we have a loss of at most 1 fiber for every backbone edge (corollary 2) and a loss of at most 4 fibers for every leg edge in each direction (inequality 1). Adding over all edges and both directions we have:

$$\begin{aligned} SOL(\mathcal{P}) &\leq OPT(\mathcal{P}) + 2 \cdot |E_{2b}| + |E_{1b}| + 4 \cdot (2 \cdot |E_{2l}| + |E_{1l}|) \leq \\ &OPT(\mathcal{P}) + 4 \cdot (2 \cdot |E_2| + |E_1|) = OPT(\mathcal{P}) + 4 \cdot |E^*| \leq 5 \cdot OPT(\mathcal{P}). \end{aligned}$$

□

5 Conclusion

In this work we studied a minimization problem occurring in multi-fiber all-optical networks with limited bandwidth, namely the MINIMUM COLLISIONS PATH MULTICOLORING problem, and presented approximation Algorithms for caterpillars, which is a commonly used network topology. We first considered the undirected version of the problem. The first of our algorithms achieves approximation ratio $1 + \frac{3|E^*|}{OPT}$ for caterpillars with only backbone paths, i.e. paths that correspond to long-distance communication requests (recall that E^* is the set of edges that are used by at least one path). The second algorithm applies to any communication pattern, and achieves approximation ratio $1 + \frac{5|E^*|}{OPT}$. Comparing to known algorithms for the more general tree topology [7] we observe that although our algorithms cannot guarantee a better approximation ratio in the worst case, they achieve much better approximation ratios for heavily loaded instances; for example, if the average load is at least $\lambda \cdot w$, for some integer λ , then it is not hard to see that $OPT \geq \lambda \cdot |E^*|$ and the above ratios become at most $1 + 3/\lambda$ and $1 + 5/\lambda$ respectively. For the directed version of the problem we

showed that our techniques can give a $1 + \frac{4|E^*|}{OPT}$ -approximation algorithm, where E^* is now the set of directed edges used by at least one path.

Our results extend to the variation with different fiber costs on Different edges, since our analysis is based on local arguments. An important observation in this context is that on backbone edges the solution is near-optimal, differing from any optimal solution by at most one fiber. This implies even better results in practice, under the reasonable assumption that the backbone fibers are more expensive to use.

An interesting open question is whether there exist better approximation algorithms, especially with respect to the worst-case ratio – the best known being the ratio 4 algorithm for trees from [7], which is based on a 4-approximation algorithm for MINCOLORS-PMC due to Chekuri, Mydlarz and Shepherd [2].

References

1. M. C. Carlisle and E. L. Lloyd. On the k -coloring of intervals. *Discrete Applied Mathematics*, 59:225–235, 1995.
2. C. Chekuri, M. Mydlarz, and F. B. Shepherd. Multicommodity demand flow in a tree. In *Proceedings 30th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 410–425, 2003.
3. R. Cole, K. Ost, and S. Schirra. Edge-coloring bipartite multigraphs in $O(E \log D)$ time. *Combinatorica*, 21(1):5–12, 2001.
4. T. Erlebach and K. Jansen. Maximizing the number of connections in optical tree networks. In *Proceedings 9th Annual International Symposium on Algorithms and Computation (ISAAC)*, volume 1533 of *LNCS*, pages 179–188. Springer, 1998.
5. T. Erlebach and K. Jansen. The complexity of path coloring and call scheduling. *Theoretical Computer Science*, 255(1–2):33–50, 2001.
6. T. Erlebach, K. Jansen, C. Kaklamanis, M. Mihail, and P. Persiano. Optimal wavelength routing in directed fiber trees. *Theoretical Computer Science*, 221:119–137, 1999.
7. T. Erlebach, A. Pagourtzis, K. Potika, and S. Stefanakos. Resource allocation problems in multifiber WDM tree networks. In *29th Workshop on Graph Theoretic Concepts in Computer Science (WG)*, Elspeet, Netherlands, 2003.
8. A. Ferreira, S. Perennes, A. Richa, H. Rivano, and N. Stier. On the design of multifiber WDM networks. In *Proceedings AlgoTel Conference*, pages 25–32, Mèze, France, 2002.
9. M. C. Golumbic and R. E. Jamison. The edge intersection graphs of paths in a tree. *Journal on Combinatorial Theory, Series B*, 38(1):8–22, 1985.
10. G. Li and R. Simha. On the wavelength assignment problem in multifiber optical tree networks. In *Terabit Optical Networking: Architecture, Control, and Management Issues*, number 4213 in *Proceedings of SPIE*, pages 84–91, 2000.
11. G. Li and R. Simha. On the wavelength assignment problem in multifiber WDM star and ring networks. *IEEE/ACM Transactions on Networking*, 9(1):60–68, 2001.
12. L. Margara and J. Simon. Wavelength assignment problem on all-optical networks with k fibres per link. In *Proceedings 27th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 768–779, 2000.
13. T. Nishizeki and K. Kashiwagi. On the 1.1 edge-coloring of multigraphs. *SIAM Journal on Discrete Mathematics*, 3(3):391–410, August 1990.
14. C. Nomikos, A. Pagourtzis, K. Potika, and S. Zachos. Path multi-coloring in weighted graphs. In *Proceedings 8th Panhellenic Conference on Informatics*, volume I, pages 178–186, Nicosia, Cyprus, 2001.

15. C. Nomikos, A. Pagourtzis, K. Potika, and S. Zachos. Fiber cost reduction and wavelength minimization in multifiber WDM networks, 2003. Submitted.
16. C. Nomikos, A. Pagourtzis, and S. Zachos. Routing and path multicoloring. *Information Processing Letters*, 80(5):249–256, 2001.
17. C. Nomikos, A. Pagourtzis, and S. Zachos. Minimizing request blocking in all-optical rings. In *Proceedings 22nd Annual Joint Conference IEEE Computer and Communications Societies (INFOCOM)*, 2003.
18. C. Nomikos, A. Pagourtzis, and S. Zachos. Satisfying a maximum number of pre-routed requests in all-optical rings. *Computer Networks: The International Journal of Computer and Telecommunications Networking*, 42(1):55–63, 2003.
19. P. Raghavan and E. Upfal. Efficient routing in all-optical networks. In *Proceedings 26th Annual ACM Symposium on Theory of Computing (STOC)*, pages 134–143, 1994.
20. P.-J. Wan and L. Liu. Maximal throughput in wavelength-routed optical networks. In *Multichannel Optical Networks: Theory and Practice*, volume 46 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 15–26. AMS, 1998.
21. P. Winkler and L. Zhang. Wavelength assignment and generalized interval graph coloring. In *Proceedings 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 830–831, Baltimore, MD, 2003.