

Scheduling Connections via Path and Edge Multicoloring^{*}

Evangelos Bampas^{1,**}, Christina Karousatou^{2,***},
Aris Pagourtzis³, and Katerina Potika⁴

¹ Univ. Bordeaux, LaBRI, UMR 5800, F-33400 Talence, France
`evangelos.bampas@labri.fr`

² LIF, Aix-Marseille University and CNRS, Marseille, France
`christina.karousatou@lif.univ-mrs.fr`

³ School of Electrical and Computer Engineering,
National Technical University of Athens, Greece
`pagour@cs.ntua.gr`

⁴ Santa Clara University
500 El Camino Real, Santa Clara, CA 95053, USA
`katerina.potika@gmail.com`

Abstract. We consider path multicoloring problems, in which one is given a collection of paths defined on a graph and is asked to color some or all of them so as to optimize certain objective functions. Typical objectives are: (a) the minimization of the average, over all edges, of the maximum-multiplicity color when the number of colors is given (MINAVGMULT-PMC), (b) the minimization of the number of colors when the maximum multiplicity for each edge is given (MIN-PMC), or (c) the maximization of the number of colored paths when both the number of colors and a maximum multiplicity constraint for each edge are given (MAX-PMC). Such problems also capture edge multicoloring variants (such as MINAVGMULT-EMC, MIN-EMC, and MAXEMC) as special cases and find numerous applications in resource allocation, most notably in optical and wireless networks, and in communication task scheduling.

Our contribution is two-fold: On the one hand, we give an exact polynomial-time algorithm for MIN-PMC on spider networks with even admissible color multiplicities on each edge. On the other hand, we present an approximation algorithm for MINAVGMULT-PMC in star networks, with a ratio strictly better than 2; our algorithm uses an appropriate path orientation. We also show that any algorithm which is based on path orientation cannot achieve an approximation ratio better than $\frac{7}{6}$.

^{*} Research supported in part by the ALGONOW project, co-financed by the European Union (European Social Fund ESF) and Greek national funds through the Operational Program Education and Lifelong Learning of the National Strategic Reference Framework (NSRF) Research Funding Program: THALES, Investing in knowledge society through the European Social Fund.

^{**} Partial support by the ANR project DISPLEXITY (ANR-11-BS02-014).

^{***} This work was done while Christina Karousatou was with the School of Electrical and Computer Engineering, National Technical University of Athens, Greece.

Our results apply to the corresponding edge multicoloring problems as well.

1 Introduction

In the class of path coloring problems, one is given a set of simple paths on a graph and is asked to assign a color to each of them so as to optimize a certain objective function. Path coloring problems have been studied extensively in the context of routing and wavelength assignment in optical networks, as well as in several other applications varying from compiler optimization to vehicle scheduling. In standard path coloring problems, edge-intersecting paths must receive distinct colors. In a meaningful generalization, path multicoloring problems were defined and studied in [18, 21, 9, 1]. Note that, by the term ‘path multicoloring’, it is meant that edge-intersecting paths may receive identical colors. This is in contrast to standard path coloring.

Various optimization objectives have been studied in the context of path multicoloring. A usual one is to minimize the number of colors used, under the assumption that, for each edge, an upper bound is given on the maximum color multiplicity allowed on that edge, i.e., on the maximum number of paths using this edge that can receive the same color. We will refer to this upper bound as the *admissible color multiplicity* for that edge. This problem is called MINIMUM PATH MULTICOLORING (MIN-PMC) and it is formally defined as follows:

Problem 1 (MINIMUM PATH MULTICOLORING, MIN-PMC).

Instance: $\langle G, \mathcal{P}, \mu \rangle$, where $G = (V, E)$ is an undirected graph, \mathcal{P} is a set of undirected simple paths on G , and $\mu : E \rightarrow \mathbb{N}$ is a function that maps each edge to the admissible color multiplicity on that edge.

Feasible solution: an assignment of one color from $\{1, \dots, w\}$ to each path in \mathcal{P} , where w is some positive integer, so that on every edge e , each color is used by at most $\mu(e)$ paths.

Goal: minimize the number of colors w .

A well-studied variant [21, 9, 25, 1] assumes a bound on the number of available colors and asks for a coloring such that the sum of edge multiplicities (equivalently, average edge multiplicity) over all edges of the graph is minimized. Here, the *multiplicity* of an edge refers to the maximum number of paths of the same color that share that edge, with respect to a certain coloring of paths. We will call this latter problem MINIMUM AVERAGE MULTIPLICITY PATH MULTICOLORING (MINAVGMULT-PMC). It is formally defined as follows:

Problem 2 (MINIMUM AVERAGE MULTIPLICITY PATH MULTICOLORING, MINAVGMULT-PMC).

Instance: $\langle G, \mathcal{P}, w \rangle$, where $G = (V, E)$ is an undirected graph, \mathcal{P} is a set of undirected simple paths on G , and $w \in \mathbb{N}$ is the number of available colors.

Feasible solution: a coloring of \mathcal{P} with w colors.

Goal: minimize $\sum_{e \in E} \mu(e)$, where $\mu(e)$ is the multiplicity of the maximum-multiplicity color on e .

Note that, in optical network terminology, MINAVGMULT-PMC corresponds to minimizing the number of parallel fiber links needed to accommodate a given set of requests with a given set of wavelengths, assuming that no wavelength collisions are allowed within the same fiber. A wavelength collision occurs when two communication requests are routed on the same fiber and modulated on the same wavelength. Other known variants include maximizing the number of satisfied requests, given constraints on both the number of colors and the admissible color multiplicities of edges (MAXIMUM PATH MULTICOLORING (MAX-PMC)), and minimizing the maximum edge multiplicity. The formal definition of MAX-PMC follows:

Problem 3 (MAXIMUM PATH MULTICOLORING, MAX-PMC).

Instance: $\langle G, \mathcal{P}, \mu, w \rangle$, where $G = (V, E)$ is an undirected graph, \mathcal{P} is a set of undirected simple paths on G , $\mu : E \rightarrow \mathbb{N}$ is a function that maps each edge to the maximum admissible color multiplicity on that edge, and $w \in \mathbb{N}$ is the number of available colors.

Feasible solution: a coloring of a subset $\mathcal{Q} \subseteq \mathcal{P}$ with w colors, so that on every edge e , each color is used by at most $\mu(e)$ paths.

Goal: maximize $|\mathcal{Q}|$.

Apart from the optical networks setting, the above problems find applications in scheduling communication requests or tasks in generic networks. In such scenarios, colors represent time slots and edge multiplicity represents congestion on the corresponding network link, assuming that requests receiving the same color are served simultaneously. Therefore, MIN-PMC is the problem of minimizing the number of time slots, given the available bandwidth per link, and MINAVGMULT-PMC corresponds to minimizing the average congestion given the number of available time slots.

Yet another interesting application comes through the known equivalence of path (multi)coloring in stars to edge (multi)coloring in general graphs. It is noteworthy that this equivalence is approximation-preserving. Therefore, one defines the corresponding MIN-EMC, MINAVGMULT-EMC and MAX-EMC problems by replacing the path intersection relation with edge incidence to the same vertex; (vertex) multiplicity now refers to the number of edges that are incident to a vertex and receive the same color. Note that MIN-EMC has been considered by Coffman et al. [15] under the name ‘ f -coloring’ as a problem of scheduling file transfers among servers that can support a certain number of parallel transfers each, in a minimum number of time slots. Drawing on their approach, we consider MINAVGMULT-EMC as a way to address the following problem: given a graph where edges represent file transfer requests, and assuming that the number of requests that can be served simultaneously by a particular node depends on the number of available ports, find a scheduling of the requests within a given number time-slots so as to minimize the total number of ports needed to satisfy all requests. Thus, solving MINAVGMULT-EMC can help a network administrator in allocating efficiently the necessary number of ports, or a network designer who has access (in advance) to demand statistics among the nodes in designing a more efficient network.

Edge (multi)-coloring problems have a self-evident connection to channel assignment problems in multi-channel wireless networks [16]. For example, different colors may represent distinct communication channels and it is possible to have parallel transmissions from the same transceiver, provided that distinct channels are used. Color multiplicity in such conditions can model the number of time frames needed to complete a given set of communication tasks. Therefore, solving MINAVGMULT-EMC amounts to minimizing the average time per node, while MIN-EMC corresponds to minimizing the number of channels needed to accomplish all tasks, given the time availability of each particular component of the network. Finally, MAX-EMC describes the situation where both the number of channels and the desired or available time per node are provided and the goal is to maximize the number of served communication requests. Under an orthogonal perspective, in a single-channel model color multiplicity may represent the available bandwidth, whereas colors may stand for time frames. Under this approach, solving MIN-EMC can help reduce the maximum time needed, MINAVGMULT-EMC minimizes the average bandwidth used, and MAX-EMC aims at maximizing the number of satisfied requests, given the bandwidth per node and the time availability. For other, more specific applications of edge multicoloring to wireless channel assignment see [14, 4] and references therein.

It makes sense to consider both undirected and directed versions of the above problems. In this paper, we provide results mainly for the undirected versions. Nonetheless, we will make use of some known results for directed versions of the problems. Throughout the text, when we refer to a particular problem, we will refer by default to its undirected version and we will state it explicitly whenever we refer to the directed version.

Related work. Path coloring and path multicoloring problems have been extensively studied during the last twenty years (see e.g. [24, 8, 7, 21, 9, 19, 20, 5, 2] and references therein). All variants are NP-hard in general graphs and even in simple topologies, e.g. stars and rings, whereas they can be solved optimally in chains. In addition, most variants are hard to approximate in general graphs within a constant factor. However, this is possible in stars and rings, and in some other simple topologies.

The MIN-PMC problem was defined independently in [9] (under the name MIN-COLORS-PMC) and in [6], where it was proved to admit a 4-approximation in trees. A $(3/2)$ -approximation for stars was proposed in [19] and an asymptotic $(9/8)$ -approximation can be obtained directly from the equivalence of MIN-PMC in stars and the MIN-EMC problem in general graphs (aka f -coloring) and the result of Coffman et al. for f -coloring [15]. In [22], algorithms for MIN-PMC in spiders and caterpillars were proposed⁵, achieving approximation ratios of 2 and 3 respectively. In contrast, it was shown in [22] that the directed version of MIN-PMC can be solved exactly in spiders.

⁵ A *spider* graph is a tree with at most one node of degree strictly greater than 2. A *caterpillar* is a tree in which all nodes of degree 2 or more lie on the same simple path.

A more recent result of Bian and Gu [5] states that MIN-PMC can be solved exactly in spiders with *uniform* and even admissible color multiplicities (here, uniform stands for ‘the same on each edge of the graph’). In addition, they show that, for stars, the uniformity requirement can be removed and, therefore, MIN-EMC can also be solved optimally in graphs where the admissible color multiplicity is even. However, Bian and Gu also prove that, in the case of odd admissible color multiplicities, the situation is not similar: the problem becomes NP-hard.

Regarding MINAVGMULT-PMC, the problem was defined in [21], and independently in [25]. In [21], an exact algorithm for chain graphs was presented, as well as 2-approximation algorithms for rings and stars. These algorithms were later extended in [20] to cover the generalized version of MINAVGMULT-PMC with non-uniform multiplicity costs. To the best of our knowledge, the MINAVGMULT-EMC problem has not been studied as such so far, hence the best known result is the 2-approximation due to the approximation preserving equivalence with MINAVGMULT-PMC in stars, as explained previously.

Results for MAX-PMC appear in [9], where a 2.54-approximation for trees is proposed, in [23], where the authors give an exact algorithm for chains and 2-approximation algorithm for rings and stars, and in [5], where they provide a 1.58-approximation algorithm for the problem in spiders and an exact algorithm for spiders with uniform and even admissible color multiplicity.

Finally, path multicoloring problems were also studied in a non-cooperative setting in [2], where the social cost depends on the maximum color multiplicity. A more general framework for studying path multicoloring under selfish considerations was proposed in [3].

Our results. In this work, we take some steps towards coping with the path and edge multicoloring problems described above.

Our first contribution is that we manage to show that MIN-PMC can be solved optimally in spiders with non-uniform even admissible color multiplicity on each edge. This represents an improvement with respect to the result of Bian and Gu [5], as it removes the uniformity requirement. As a byproduct, we obtain an exact algorithm for MAX-PMC in such graphs, i.e., spiders with non-uniform even color multiplicity. This latter result holds for any number of given colors.

We then turn our focus to MINAVGMULT-EMC (equivalent to MINAVGMULT-PMC in stars) and start by rewriting an algorithm of [21] in edge multicoloring terms. This algorithm, which achieves an approximation factor of 2, has a locally near-optimal behaviour: the maximum multiplicity at each node differs from the optimal by at most one. We take advantage of this property and first show that MINAVGMULT-EMC can be solved exactly if the number of available colors equals 2. Next, we further fine-tune the algorithm, by adding a random orientation step together with a derandomization process, in order to come up with an algorithm having an approximation ratio of $2 - \frac{1}{2^w}$, where w is the number of available colors. Since the ratio of 2 of the algorithm in [21] is tight, our results represent the first, to the best of our knowledge, algorithm with an approxima-

tion ratio strictly better than 2, for both MINAVGMULT-EMC in general graphs and MINAVGMULT-PMC in stars.

We conclude by proving a $\frac{7}{6}$ lower bound for all orientation-based algorithms for MINAVGMULT-EMC, i.e., for algorithms that attempt to solve the (unoriented) edge (or path) multicoloring problem by computing an orientation of the given multigraph (or set of paths) and then applying the algorithm from [21].

Remark 1 (Notation). Throughout the paper, we will use the notation L_e for the load of edge e with respect to a given path set, i.e., the number of paths which contain edge e . Also, d_v will denote the degree of a node v in the context of an undirected graph, whereas for directed graphs we will use d_v^{in} (resp. d_v^{out}) for the in-degree (resp. out-degree) of node v .

Some of the proofs are omitted due to lack of space and they are deferred to the full version of the paper.

2 Min-PMC and Max-PMC on Spiders

In this section, we present exact algorithms for MIN-PMC and MAX-PMC on spiders with non-uniform even admissible color multiplicity.

2.1 Minimizing the number of colors

With respect to the MIN-PMC problem in any graph, one observes immediately the following lower bound:

Fact 1. *No MIN-PMC instance can be colored with fewer than $w_{\text{lb}} = \max_e \left\lceil \frac{L_e}{\mu(e)} \right\rceil$ colors.*

It is known [24, 8] that MIN-PMC is NP-hard when G is a star and $\mu(e) = 1$ for all edges. Bian and Gu [5] prove that for every odd k , MIN-PMC is NP-hard when G is a star and $\mu(e) = k$ for all edges. They also give a polynomial-time algorithm for the case where G is a spider and every leg of the spider has uniform and even admissible color multiplicity. We extend their algorithm in order to remove the uniformity requirement and we show that MIN-PMC is polynomial-time solvable on spiders with even admissible color multiplicity.

We reduce the problem to the *directed* version of MIN-PMC, which is known to be polynomial-time solvable on spiders [22]. Formally, the directed version of the problem is defined as follows:

Problem 4 (DIRECTED MINIMUM PATH MULTICOLORING, DIR-MIN-PMC).

Instance: $\langle G, \vec{\mathcal{P}}, \mu \rangle$, where $G = (V, E)$ is an undirected graph, $\vec{\mathcal{P}}$ is a set of directed simple paths on G , and $\mu : E \rightarrow \mathbb{N}$ is a function that maps each edge to the admissible color multiplicity on that edge for each direction.

Feasible solution: a coloring of $\vec{\mathcal{P}}$ so that on every edge e , each color is used by at most $\mu(e)$ paths in each direction.

Goal: minimize the number of colors.




Fig. 1. Original MIN-PMC instance (left) and instance with added dummy paths (right).




Fig. 2. The corresponding graph H (left) and an orientation of the closed paths in its Euler partition (right).




Fig. 3. The DIR-MIN-PMC instance.

In order to perform the reduction, we need to decide on a direction for each undirected path in the original MIN-PMC instance $\langle G, \mathcal{P}, \mu \rangle$. We accomplish this by adding unit-length dummy paths on edges with odd load, and then considering the multigraph H which has the same node set as G and contains one edge (u, v) for each path with endpoints u and v . It can be shown that the addition of dummy paths ensures that all nodes of H have even degree. Therefore, the Euler partition algorithm of [10] will partition the edges of G into closed paths. We then perform an arbitrary orientation of each closed path in the Euler partition. This assigns a direction to each edge of H , which in turn corresponds to a direction for each path in \mathcal{P} and for each dummy path. Figures 1–3 illustrate the reduction.

Theorem 1. *Algorithm 1 is an exact polynomial-time algorithm for MIN-PMC on spiders with even admissible color multiplicity.*

Algorithm 1 An exact algorithm for MIN-PMC in spiders with even admissible color multiplicity

Input: an instance $\langle G, \mathcal{P}, \mu \rangle$ of MIN-PMC, where $G = (V, E)$ is a spider and $\mu(e)$ is even for all e

Output: an optimal solution using w_{1b} colors

- 1: For any edge with odd load, add one unit-length dummy path on that edge. Let \mathcal{P}' be the resulting set of paths.
 - 2: Construct the multigraph $H = (V, E')$, where E' contains one edge (u, v) for each path in \mathcal{P}' with endpoints u and v .
 - 3: Use the algorithm EP in [10] to find an Euler partition of H . Orient each closed path of the Euler partition in an arbitrary manner, thus obtaining a directed multigraph \vec{H} . Assign a direction to each path in \mathcal{P}' according to the orientation of the corresponding edge in \vec{H} , thus obtaining a set $\vec{\mathcal{P}}'$ of directed paths on G .
 - 4: Find an optimal coloring c of the DIR-MIN-PMC instance $\langle G, \vec{\mathcal{P}}', \mu' \rangle$, where $\mu'(e) = \frac{\mu(e)}{2}$ for all e , using the algorithm in [22, Theorem 6].
 - 5: Return the coloring c restricted to the original paths in \mathcal{P} .
-

2.2 Maximizing the number of satisfied requests

A corollary of Theorem 1 is that the MAX-PMC problem is also optimally solvable in polynomial time on spiders with (non-uniform) even admissible color multiplicity.

In [5], Bian and Gu propose a deterministic polynomial-time algorithm for MAX-PMC on spiders with uniform even admissible color multiplicity, which works as follows:

1. First, it computes a maximum-size subset \mathcal{Q} of \mathcal{P} such that every edge e of the spider is used by at most $w\mu(e)$ paths in \mathcal{Q} . This is accomplished by solving optimally an instance of the Call Control problem on spiders, via an algorithm developed in [5].
2. Then, \mathcal{Q} is colored optimally with w colors using the algorithm for MIN-PMC on spiders with uniform even admissible color multiplicity.

By plugging in Algorithm 1 instead in step 2, we can solve MAX-PMC on spiders with non-uniform admissible color multiplicity. Thus, we have the following:

Theorem 2. *There exists an exact polynomial-time algorithm for MAX-PMC on spiders with even admissible color multiplicity.*

3 Reducing Average Color Multiplicity on Stars

In this section, we present our results on the approximability of MINAVGMULT-PMC on stars. We will only be interested in MINAVGMULT-PMC instances in which all paths in \mathcal{P} are of length 2. In view of the following proposition, we can essentially ignore paths of length 1 if we ensure that our algorithm produces solutions whose cost is at most a constant times $\text{OPT}_{LB}(\mathcal{I}) = \sum_{e \in E} \lceil \frac{L_e}{w} \rceil$ (a lower bound on the cost of any optimal solution of \mathcal{I}):

Proposition 1. *Let \mathcal{A} be a polynomial-time algorithm for MINAVGMULT-PMC instances with paths of length 2 which, for any such instance \mathcal{I} , produces a solution with cost $\text{SOL}_{\mathcal{A}}(\mathcal{I}) \leq \alpha \cdot \text{OPT}_{LB}(\mathcal{I})$, where $\alpha \geq 1$ is a constant. Then, there exists a polynomial-time algorithm \mathcal{B} with the same approximation guarantee for general instances \mathcal{I}' of MINAVGMULT-PMC.*

It has been pointed out before in the literature [11, 8] that there exists an easily computed bijection between pairs (G, \mathcal{P}) , where G is a star and \mathcal{P} is a set of paths of length 2 on G , and multigraphs G' , such that the conflict graph of \mathcal{P} is isomorphic to the line graph of G' . Therefore, MINAVGMULT-PMC can be equivalently cast as the following problem:

Problem 5 (MINIMUM AVERAGE MULTIPLICITY EDGE MULTICOLORING, MINAVGMULT-EMC).

Instance: $\langle G, w \rangle$, where $G = (V, E)$ is an undirected multigraph and $w \in \mathbb{N}$ is the number of available colors.

Feasible solution: a coloring of E with w colors.

Goal: minimize $\sum_{v \in V} \mu(v)$, where $\mu(v)$ is the multiplicity of the maximum-multiplicity color over all edges incident to v .

Fact 2. *The multiplicity of the maximum multiplicity color incident to node v is at least $\lceil \frac{d_v}{w} \rceil$, thus the minimum cost for any MINAVGMULT-EMC instance is at least $\sum_{v \in V} \lceil \frac{d_v}{w} \rceil$.*

For any fixed $w \geq 3$, MINAVGMULT-EMC is NP-hard via a straightforward reduction from the decision version of the classical edge coloring problem on w -regular graphs, which is known to be NP-complete [13, 17]. Nomikos et al. [21] propose a 2-approximation algorithm which we restate as Algorithm 2 in MINAVGMULT-EMC terms (the algorithm was originally stated in MINAVGMULT-PMC terms). The analysis in [21] is tight, as there exists a family of instances on which Algorithm 2 computes a solution with cost exactly twice the optimum: $\{\langle C_k, w \rangle : k \geq 2 \text{ and } w \geq 2\}$, where C_k is the ring graph with k nodes. Indeed, if the directions assigned in step 1 are such that each node has in-degree 1 and out-degree 1, then the resulting solution will have cost $2k$, whereas the optimum solution has cost k .

If we have only two available colors, the problem can be solved exactly in polynomial time: The Euler partition algorithm in [10] computes a partition of the edges of a multigraph into open and closed paths, with the property that each vertex of odd degree is the end of exactly one open path, and each vertex of even degree is the end of no open paths. Note, then, that if we color the edges of each path of the Euler partition alternately with the two available colors, the resulting coloring will have the property that the edges incident to each even-degree node will be equipartitioned into two same-colored subsets, whereas the edges incident to each odd-degree node will be partitioned into two same-colored subsets whose sizes differ by exactly one. This implies that the cost incurred by each node v will be exactly $\lceil \frac{d_v}{2} \rceil$, thus the solution is optimal in view of Fact 2. We have thus proved the following:

Algorithm 2 A 2-approximation algorithm for MINAVGMULT-EMC [21]

Input: an instance $\langle G, w \rangle$ of MINAVGMULT-EMC, $G = (V, E)$

Output: a 2-approximate solution

- 1: Assign an arbitrary direction to each edge of G .
 - 2: For each $v \in V$, group its out_v outgoing edges into $\lceil \frac{\text{out}_v}{w} \rceil$ groups of at most w edges each, and let V_{out} denote the set of all groups of outgoing edges. Similarly, for each $v \in V$, group its in_v incoming edges into $\lceil \frac{\text{in}_v}{w} \rceil$ groups of at most w edges each, and let V_{in} denote the set of all groups of incoming edges.
 - 3: Construct the bipartite multigraph $H = (V_{\text{out}} \cup V_{\text{in}}, A)$, where for each edge in E , A contains one edge joining its outgoing group to its incoming group. The maximum degree of H is bounded by w .
 - 4: Use an algorithm for edge coloring of bipartite multigraphs (e.g. [10]) to compute an edge coloring of H with w colors.
 - 5: Assign to each edge of G the color of the corresponding edge in H .
-

Theorem 3. *There exists an exact polynomial-time algorithm for MINAVGMULT-EMC with two available colors.*

3.1 An approximation algorithm with ratio strictly better than 2

We first show that if we assign random directions to the edges of a multigraph with n nodes, then the resulting orientation will have, in expectation, at least $\frac{1}{2w} \cdot n$ nodes such that, if we subsequently execute steps 2-5 of Algorithm 2 on this orientation, each of them will contribute the minimum possible cost to the cost of the solution: $\lceil \frac{d_v}{w} \rceil$. Then, we show how to derandomize this procedure in order to obtain a deterministic algorithm for MINAVGMULT-EMC with approximation ratio at most $(2 - \frac{1}{2w})$.

As a preliminary observation, consider the simplest conceivable randomized algorithm for MINAVGMULT-EMC, i.e. the one that assigns a random color with uniform probability to each edge. Unfortunately, this algorithm performs quite poorly in the following family of instances: The multigraph contains two nodes with w parallel edges joining them, where w is the number of available colors. As is well known [12], the maximum multiplicity color will appear $\Theta\left(\frac{\log w}{\log \log w}\right)$ times with high probability, whereas the optimum cost is 2. This motivates us to search for a randomized algorithm with a better performance guarantee.

Definition 1 (Locally optimal nodes). *Let $\langle G, w \rangle$ be an instance of the MINAVGMULT-EMC problem, let \vec{G} be an orientation of G , and let v be a node of G . We say that v is locally optimal if the following condition holds:*

$$(d_v^{\text{in}} \bmod w = 0) \vee (d_v^{\text{out}} \bmod w = 0) \vee ((d_v^{\text{in}} \bmod w) + (d_v^{\text{out}} \bmod w) > w)$$

The following lemma is implicit in the analysis in [21].

Lemma 1 ([21]). *In any solution computed by Algorithm 2, each node v incurs a cost of exactly $\lceil \frac{d_v}{w} \rceil$ if it is locally optimal with respect to the directions assigned after step 1, or $\lceil \frac{d_v}{w} \rceil + 1$ if it is not locally optimal.*

In other words, Algorithm 2 incurs an additional cost, with respect to the lower bound of Fact 2, of exactly one for each non-locally-optimal node. The existence of an algorithm that computed an orientation of G in which at least some fraction of the nodes were guaranteed to be locally optimal would yield an approximation algorithm for MINAVGMULT-EMC with ratio strictly better than 2. We now consider the algorithm that assigns a random direction to each edge of G and then executes steps 2-5 of Algorithm 2.

Definition 2. $S(d, w)$ denotes the set of integers i such that, if exactly i out of d incident edges to a d -degree node are incoming and $(d - i)$ incident edges are outgoing, then the node is locally optimal assuming we have w colors.

Lemma 2. If d is a multiple of w , then $S(d, w) = \{i \cdot w : 0 \leq i \leq \frac{d}{w}\}$.

Lemma 3. If $d = r \cdot w + x$, where $0 < x < w$, then $S(d, w)$ is exactly the set: $\bigcup_{i=0}^r \{i \cdot w + j : j \in \{0, x, x + 1, \dots, w - 1\}\} \cap \{0, 1, \dots, d\}$.

Lemma 4. Let $\langle G, w \rangle$ be an instance of MINAVGMULT-EMC and let \vec{G} be a random orientation of G in which each edge receives each of the two possible directions with probability $\frac{1}{2}$. The expected fraction of locally optimal nodes in \vec{G} is at least $\frac{1}{2^w}$.

Theorem 4. There exists a $(2 - \frac{1}{2^w})$ -approximation algorithm for MINAVGMULT-EMC.

Proof. By Lemma 4, if we assign random directions to the edges of G , we get at least $\frac{1}{2^w} \cdot n$ locally optimal nodes in expectation. This algorithm can be derandomized by a straightforward application of the method of conditional expectations. Indeed, if we assume that the orientation of a subset of the edges has already been fixed, we can compute in polynomial time the probability that a fixed node v of degree d will be locally optimal if we assign the rest of the directions randomly, as follows: Assume that a of its incident edges have already been oriented as incoming to v , and b of its incident edges have already been oriented as outgoing from v . Then, the probability that the node will be locally optimal is $\frac{1}{2^{d-a-b}} \cdot \sum_{s: s+a \in S(d, w)} \binom{d-a-b}{s}$.

Therefore, the algorithm which examines edges in an arbitrary order, and to each edge assigns the direction which maximizes the expected number of locally optimal nodes under the current partial orientation, runs in deterministic polynomial time and produces an orientation with at least $\frac{1}{2^w} \cdot n$ locally optimal nodes. Taking also into account Lemma 1, this implies that, if we execute steps 2-5 of Algorithm 2 on this orientation, we will obtain a solution with cost SOL that can be expressed as follows: (let V denote the set of nodes of G and \mathcal{O}




Fig. 4. The graphs G_1 , G_2 , and G_3 , as defined in Section 3.2.

denote the set of locally optimal nodes)

$$\begin{aligned}
\text{SOL} &= \sum_{v \in \mathcal{O}} \left\lfloor \frac{d_v}{w} \right\rfloor + \sum_{v \in V \setminus \mathcal{O}} \left(\left\lfloor \frac{d_v}{w} \right\rfloor + 1 \right) \\
&= \sum_{v \in V} \left\lfloor \frac{d_v}{w} \right\rfloor + |V \setminus \mathcal{O}| \\
&\leq \sum_{v \in V} \left\lfloor \frac{d_v}{w} \right\rfloor + \left(1 - \frac{1}{2^w} \right) \cdot |V|. \tag{1}
\end{aligned}$$

If OPT is the cost of an optimal solution, then Eq. 1, Fact 2, and the observation that $\text{OPT} \geq |V|$ yield: $\text{SOL} \leq \left(2 - \frac{1}{2^w} \right) \cdot \text{OPT}$. \square

3.2 A lower bound for orientation-based algorithms

Let $\{G_i\}_{i \geq 1}$ be the infinite family of graphs where $G_1 = K_4$, and G_{i+1} is constructed by attaching to each degree-3 node of G_i a new clique K_4 (by identifying the degree-3 node of G_i to which the K_4 is attached with one of the nodes of the attached clique). Figure 4 illustrates the first few graphs in the family. For every $i \geq 1$, consider the instance $\mathcal{I}_i = \langle G_i, 3 \rangle$ of MINAVGMULT-EMC .

A simple calculation reveals that the number of nodes n_i in G_i is $n_i = 2 \cdot (3^i - 1)$. In particular, for $i \geq 2$, the number of nodes of degree 3 in G_i is $n'_i = 4 \cdot 3^{i-1}$, whereas the rest of the nodes have degree 6.

For every $i \geq 2$, there exists a coloring of the edges of G_i , such that each node v contributes exactly $\left\lfloor \frac{d_v}{w} \right\rfloor$ to the cost of the solution (Figure 5). By Fact 2, this coloring is necessarily optimal. Therefore, the cost of the optimal solution for \mathcal{I}_i is: $\text{OPT}_i = n'_i + 2(n_i - n'_i) = \frac{8}{3} \cdot 3^i - 4$.

On the other hand, for $i \geq 2$, for every orientation of G_i , out of every three nodes of degree 3 which belong to the same K_4 subgraph of G_i , at most two will be locally optimal. Indeed, these nodes form a triangle in G_i , and thus any orientation of the edges between them will result in at least one of the nodes




Fig. 5. An optimal coloring of K_4 with three colors. We obtain an optimal coloring for \mathcal{I}_i by repeating this coloring for every K_4 clique contained in G_i .

having at least one outgoing and at least one incoming edge. Therefore, every MINAVGMULT-EMC algorithm which computes an orientation of the given multigraph and then executes steps 2-5 of Algorithm 2, will compute a solution for \mathcal{I}_i with cost: $\text{SOL}_i \geq 2 \cdot \frac{n'_i}{3} + \frac{2n'_i}{3} + 2 \cdot (n_i - n'_i) = \frac{28}{9} \cdot 3^i - 4$.

By the calculations for OPT_i and SOL_i and letting $i \rightarrow +\infty$, we obtain:

Theorem 5. *No MINAVGMULT-EMC algorithm which computes an orientation of the given multigraph and then executes steps 2-5 of Algorithm 2 can have an approximation ratio which is better than $\frac{7}{6}$.*

4 Concluding Remarks

We gave exact algorithms for MIN-PMC and MAX-PMC on spiders with (non-uniform) even admissible color multiplicity, thus extending the algorithms from [5] which work only for spiders with uniform even admissible color multiplicity.

Furthermore, we gave an approximation algorithm for MINAVGMULT-PMC on stars with approximation ratio at most $(2 - \frac{1}{2w})$. This improves a previous algorithm from [21], which guarantees an approximation ratio of 2. Note that the bound of 2 on the approximation ratio of the algorithm from [21] is tight. Having said that, one could easily modify that algorithm to guarantee an approximation ratio of $(2 - \frac{1}{|E|})$, where $|E|$ is the number of edges of the star network. However, from a practical standpoint, in most applications the number of colors w represents a scarce network resource which needs to be utilized efficiently (e.g., in particular in the context of optical networks, w represents the number of available wavelengths per fiber, which is very limited due to technological constraints), whereas the size of the network can grow unbounded. Therefore, the approximation ratio of our algorithm represents a significant improvement over the previous algorithms.

Finally, we have provided a lower bound of $\frac{7}{6}$ on the approximation ratio of any algorithm which computes an orientation of the edges of the given multigraph and then executes steps 2-5 of Algorithm 2. In view of the large gap between this lower bound and the upper bound of Theorem 4, it makes sense to try to improve the performance of the edge orientation algorithm. For example, one could try to tweak the probabilities with which edges receive one of the two directions, in order to increase the expected number of locally optimal nodes.

References

1. Andrews, M., Zhang, L.: Minimizing maximum fiber requirement in optical networks. *J. Comput. Syst. Sci.* 72(1), 118–131 (2006), <http://dx.doi.org/10.1016/j.jcss.2005.08.001>
2. Bampas, E., Pagourtzis, A., Pierrakos, G., Potika, K.: On a noncooperative model for wavelength assignment in multifiber optical networks. *IEEE/ACM Trans. Netw.* 20(4), 1125–1137 (2012), <http://dx.doi.org/10.1109/TNET.2011.2173948>
3. Bampas, E., Pagourtzis, A., Pierrakos, G., Syrgkanis, V.: Selfish resource allocation in optical networks. In: Spirakis, P.G., Serna, M.J. (eds.) *Algorithms and Complexity*, 8th International Conference, CIAC 2013, Barcelona, Spain, May 22–24, 2013. Proceedings. *Lecture Notes in Computer Science*, vol. 7878, pp. 25–36. Springer (2013), <http://dx.doi.org/10.1007/978-3-642-38233-8>
4. Barrett, C.L., Kumar, V.S.A., Marathe, M.V., Thite, S., Istrate, G.: Strong edge coloring for channel assignment in wireless radio networks. In: 4th IEEE Conference on Pervasive Computing and Communications Workshops (PerCom 2006 Workshops), 13–17 March 2006, Pisa, Italy. pp. 106–110. IEEE Computer Society (2006), <http://doi.ieeecomputersociety.org/10.1109/PERCOMW.2006.129>
5. Bian, Z., Gu, Q.P.: Wavelength assignment in multifiber star networks. *Networks* 56(1), 30–38 (2010)
6. Chekuri, C., Mydlarz, M., Shepherd, F.B.: Multicommodity demand flow in a tree. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) *Automata, Languages and Programming*, 30th International Colloquium, ICALP 2003, Eindhoven, The Netherlands, June 30 - July 4, 2003. Proceedings. *Lecture Notes in Computer Science*, vol. 2719, pp. 410–425. Springer (2003)
7. Erlebach, T., Jansen, K.: Maximizing the number of connections in optical tree networks. In: Chwa, K., Ibarra, O.H. (eds.) *Algorithms and Computation*, 9th International Symposium, ISAAC '98, Taejon, Korea, December 14–16, 1998, Proceedings. *Lecture Notes in Computer Science*, vol. 1533, pp. 179–188. Springer (1998)
8. Erlebach, T., Jansen, K.: The complexity of path coloring and call scheduling. *Theoretical Computer Science* 255(1–2), 33–50 (2001)
9. Erlebach, T., Pagourtzis, A., Potika, K., Stefanakos, S.: Resource allocation problems in multifiber WDM tree networks. In: Bodlaender, H.L. (ed.) *Graph-Theoretic Concepts in Computer Science*, 29th International Workshop, WG 2003, Elspeet, The Netherlands, June 19–21, 2003, Revised Papers. *Lecture Notes in Computer Science*, vol. 2880, pp. 218–229. Springer (2003)
10. Gabow, H.N.: Using Euler partitions to edge color bipartite multigraphs. *International Journal of Computer and Information Sciences* 5(4), 345–355 (1976)
11. Golumbic, M.C., Jamison, R.E.: The edge intersection graphs of paths in a tree. *Journal of Combinatorial Theory, Series B* 38(1), 8–22 (1985)
12. Gonnet, G.H.: Expected length of the longest probe sequence in hash code searching. *Journal of the ACM* 28(2), 289–304 (1981)
13. Holyer, I.: The NP-completeness of edge-colouring. *SIAM Journal on Computing* 10(4), 718–720 (1981)
14. Hsu, C., Liu, P., Wang, D., Wu, J.: Generalized edge coloring for channel assignment in wireless networks. In: 2006 International Conference on Parallel Processing (ICPP 2006), 14–18 August 2006, Columbus, Ohio, USA. pp. 82–92. IEEE Computer Society (2006), <http://dx.doi.org/10.1109/ICPP.2006.45>

15. Jr., E.G.C., Garey, M.R., Johnson, D.S., LaPaugh, A.S.: Scheduling file transfers. *SIAM J. Comput.* 14(3), 744–780 (1985), <http://dx.doi.org/10.1137/0214054>
16. Kyasanur, P., Vaidya, N.H.: Capacity of multichannel wireless networks under the protocol model. *IEEE/ACM Trans. Netw.* 17(2), 515–527 (2009), <http://doi.acm.org/10.1145/1552193.1552205>
17. Leven, D., Galil, Z.: NP completeness of finding the chromatic index of regular graphs. *Journal of Algorithms* 4(1), 35–44 (1983)
18. Li, G., Simha, R.: On the wavelength assignment problem in multifiber WDM star and ring networks. In: *INFOCOM*. pp. 1771–1780 (2000), <http://www.ieee-infocom.org/2000/papers/222.ps>
19. Nomikos, C., Pagourtzis, A., Potika, K., Zachos, S.: Fiber cost reduction and wavelength minimization in multifiber WDM networks. In: Mitrou, N., Kontovasilis, K.P., Rouskas, G.N., Iliadis, I., Merakos, L.F. (eds.) *NETWORKING 2004, Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communication*, Third International IFIP-TC6 Networking Conference, Athens, Greece, May 9-14, 2004, Proceedings. *Lecture Notes in Computer Science*, vol. 3042, pp. 150–161. Springer (2004)
20. Nomikos, C., Pagourtzis, A., Potika, K., Zachos, S.: Routing and wavelength assignment in multifiber WDM networks with non-uniform fiber cost. *Computer Networks* 50(1), 1–14 (2006), <http://dx.doi.org/10.1016/j.comnet.2004.11.028>
21. Nomikos, C., Pagourtzis, A., Zachos, S.: Routing and path multicoloring. *Information Processing Letters* 80(5), 249–256 (2001)
22. Pagourtzis, A., Potika, K., Zachos, S.: Path multicoloring with fewer colors in spiders and caterpillars. *Computing* 80(3), 255–274 (2007)
23. Potika, K.: Maximizing the number of connections in multifiber WDM chain, ring and star networks. In: Boutaba, R., Almeroth, K.C., Puigjaner, R., Shen, S.X., Black, J.P. (eds.) *NETWORKING 2005: Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communication Systems*, 4th International IFIP-TC6 Networking Conference, Waterloo, Canada, May 2-6, 2005, Proceedings. *Lecture Notes in Computer Science*, vol. 3462, pp. 1465–1470. Springer (2005)
24. Raghavan, P., Upfal, E.: Efficient routing in all-optical networks. In: Leighton, F.T., Goodrich, M.T. (eds.) *STOC*. pp. 134–143. ACM (1994)
25. Winkler, P., Zhang, L.: Wavelength assignment and generalized interval graph coloring. In: *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, January 12-14, 2003, Baltimore, Maryland, USA. pp. 830–831. ACM/SIAM (2003), <http://dl.acm.org/citation.cfm?id=644108.644246>