Profit Maximization in Flex-Grid All-Optical Networks

Mordechai Shalom[2], Prudence W.H. Wong[1], Shmuel Zaks[3]

[1] Department of Computer Science, University of Liverpool, UK
pwong@liverpool.ac.uk
[2] TelHai Academic College, Upper Galilee, 12210, Israel, cmshalom@telhai.ac.il
[3] Department of Computer Science, Technion, Haifa, Israel,
zaks@cs.technion.ac.il. Visiting Professor, Departamento de Ingenieria Telematica, Universidad Carlos III de Madrid, Spain, and Visiting Researcher, Institute IMDEA Networks, Madrid, Spain.

Abstract. All-optical networks have been largely investigated due to their high data transmission rates. The key to the high speeds in all-optical networks is to maintain the signal in optical form, to avoid the overhead of conversion to and from electrical form at the intermediate nodes. In the traditional WDM technology the spectrum of light that can be transmitted through the optical fiber has been divided into frequency intervals of fixed width with a gap of unused frequencies between them. In this context the term wavelength refers to each of these predefined frequency intervals.
An alternative architecture emerging in very recent studies is to move towards a flexible model in which the usable frequency intervals are of variable width. Every lightpath is assigned a frequency interval which remains fixed through all the links it traverses. Two different lightpaths using the same link have to be assigned disjoint sub-spectra. This technology is termed *flex-grid* or *flex-spectrum*.
The introduction of this technology requires the generalization of many optimization problems that have been studied for the fixed-grid technology. Moreover it implies new problems that are irrelevant or trivial in the current technology. In this work we focus on bandwidth utilization in path toplogy and consider two *wavelength assignment*, or in graph theoretic terms *coloring*, problems where the goal is to maximize the total profit. We obtain bandwidth maximization as a special case.

Keywords: all-optical networks, flex-grid, approximation algorithms, network design.

1 Introduction

The WDM technology: All-optical networks have been largely investigated in recent years due to the promise of high data transmission rates. Its major applications are in video conferencing, scientific visualization, real-time medical imaging, high-speed super-computing, cloud computing, distributed computing, and media-on-demand. The key to high speeds in all-optical networks is to maintain the signal in optical form, thereby avoiding the prohibitive overhead of conversion to and from the electrical form at the intermediate nodes.

In modern optical networks, high-speed signals are sent through optical fibers using WDM (Wavelength Division Multiplexing) technology: several signals connecting different source - destination pairs may share a link, provided they are transmitted on carriers having different wavelengths of light. These signals are routed at intermediate nodes by optical cross-connects (OXCs) that can route an incoming signal arriving from an incident edge to another, based on the signal's wavelength. A signal transmitted optically from some source node to some destination node over a wavelength is termed a *lightpath*.

Fixed-grid and flex-grid DWDM networks: Traditionally the spectrum of light that can be transmitted through the fiber has been divided into frequency intervals of fixed width with a gap of unused frequencies between them. In this context the term wavelength refers to each of these predefined frequency intervals. This technology is termed WDM, DWDM or UDWDM depending on the gap of unused frequencies between the wavelengths.

An alternative architecture emerging in very recent studies is to move away from this rigid DWDM model towards a flexible model in which the usable frequency intervals are of variable width (even within the same link). Every lightpath has to be assigned a frequency interval (*sub-spectrum*), which remains fixed through all the links it traverses. As in the traditional model, two different lightpaths using the same link have to be assigned disjoint sub-spectra. This technology is termed *flex-grid* or *flex-spectrum*, as opposed to *fixed-grid* or *fixed-spectrum* current technology. Specifically this new technology is feasible due to gridless wavelength selective switches (WSS), based on a very large number of pixels. This *sliceable* transceiver technology is not as mature, but is critical to the economic viability of flex-grid.

The introduction of the flex-grid technology requires the generalization of most of the many optimization problems that have been studied under the fixed-grid technology. For instance, as a result of the variability of the width of the sub-spectra, lightpaths have different transmission impairments, thus different regeneration needs. Another major difference is that in the fixed-grid it is assumed that lightpath requests are for one wavelength's bandwidth because otherwise it can be treated as multiple independent requests. In the flex-grid technology this assumption does not hold because two lightpaths assigned two arbitrary wavelengths are not equivalent to one lightpath assigned two consecutive colors. This assignments differ both in terms of regeneration needs, and in terms of bandwidth utilization.

In this work we focus on the bandwidth utilization in path topology as a basic network to analyze in this introductory work. Results on path topology may extend to rings and trees that are other natural topologies in optical networks. Such results often have applications in the scheduling context in which the path network becomes the time axis. For problems that are provably hard in the general case we consider special cases such as bounded load and proper intervals.

We assume that the lightpath requests have bandwidth requirements that are multiples of some basic unit. This unit is smaller than the traditional wavelength bandwidth. The entire bandwidth of the fiber is W units. We consider two *wavelength assignment*, or in graph theoretic terms *coloring*, problems. In both problems every lightpath request consists of a path P, with minimum and maximum bandwidth requirements $\mathbf{a}_P$ and $\mathbf{b}_P$ respectively, and a unit profit $\mathbf{w}_P$ (i.e., the profit for each color assigned). In the first problem such a lightpath P has to be assigned a set $w(P)$ of colors such that $\mathbf{a}_p \leq |w(P)| \leq \mathbf{b}_P$ where color is a number between 0 and $W-1$. In the second problem, in addition, the set $w(P)$ of colors assigned to a lightpath P constitutes an interval of colors from some color λ to some color $\lambda' \geq \lambda$ so that the loss, due to the otherwise unused gap between the colors, is avoided. We term these colorings as *non-contiguous colorings* (or just *colorings*), and *contiguous colorings* respectively. Note that these colorings correspond to ordinary colorings and to *interval colorings* of the intersection graph of the paths.

The profit obtained from a lightpath is proportional to the number of colors it is assigned and its unit profit, i.e. $\mathbf{w}_P \cdot |w(P)|$. Our goal is to maximize the total profit. We have an important special case when $\mathbf{w}_P$ is equal to the length of the path P. In this case the profit is the total bandwidth utilization of the network.

Related Work: [1] is a general reference for optical networks. For a discussion of their data transmission rates see [2]. [3, 4] suggest flex-grid DWDM as an alternative emerging architecture. The network implications of this new architecture are explained in detail in [5], which refers to the key enabling technologies for the flex-spectrum.

Closely related to our work is the coloring and interval coloring of interval graphs. [6] is an excellent reference book on these subjects. To find an interval coloring with minimum colors in an interval graph is known as the *shipbuilding* problem, and also as the *dynamic storage allocation* problem. The problem is stated in [7] as NP-Complete under the latter name (problem [SR2]). In [8] it is conjectured to be in APX-Hard. Interval coloring of interval graphs with different optimization functions have also been studied in the literature (see for instance [9]).

Our Contribution: In this paper we consider three profit maximization problems, PMC is for non-contiguous coloring, PMCC is for contiguous coloring and PMCCC is for circularly contiguous coloring. Circularly contiguous coloring means that the interval of colors assigned can be wrapped around from $W-1$ back to 0. For PMC, we show a polynomial time optimal algorithm for arbitrary $\mathbf{a}$ and $\mathbf{b}$ when the network is a path. For PMCC, we derive an algorithm that converts a circularly contiguous coloring to a contiguous coloring with a small loss in the profit. We observe that PMCC is NP-Hard for path networks and study special cases. We study the case when the number of paths that using any given edge is bounded by some constant and give a polynomial time optimal algorithm. We further consider the case when the input set of paths is proper, i.e., no path properly contains another, and show an approximation algorithm with approximation ratio 4/3 for some special values of **a** and **b**.

2 Preliminaries

Graphs and paths: In path (multi) coloring problems we are given a network modeled by a graph G and a set of lightpaths modeled by a set $\mathcal{P}$ of non-trivial paths of G. $V(G)$ and $E(G)$ denote the vertex set and edge set of G, respectively. We denote by $\delta_G(v)$ the set of edges incident to a vertex v in G, i.e. $\delta_G(v) = \{e \in E(G) | v \in e\}$, and $d_G(v) = |\delta_G(v)|$ is the degree of v in G. For a directed graph G, $A(G)$ denotes the arc set of G. We denote by $\delta^-_G(v)$ and $\delta^+_G(v)$ the sets of incoming arcs and outgoing arcs of a vertex v, respectively. Similarly $d^-_G(v) = \left|\delta^-_G(v)\right|$ (resp. $d^+_G(v) = \left|\delta^+_G(v)\right|$) denotes the in-degree (resp. out-degree) of v in G.

We consider paths as sets of edges, e.g. for two paths P, P' we denote by $P \cap P'$ the set of their common edges, and by $|P|$ the length of P. For an edge e of G, we denote by $\mathcal{P}_e$ the subset of $\mathcal{P}$ consisting of the paths containing e, i.e. $\mathcal{P}_e \stackrel{def}{=} \{P \in \mathcal{P} : e \in P\}$. The number of these paths is termed the *load* on the edge e, and denoted by $L_e(\mathcal{P}) \stackrel{def}{=} |\mathcal{P}_e|$. An important parameter we consider is the maximum load over all the edges of G. We denote it by $L_{max}(\mathcal{P}) \stackrel{def}{=} \max\{L_e(\mathcal{P}) : e \in E(G)\}$. Note that in the intersection graph of the paths $\mathcal{P}$, the subset of vertices corresponding to $\mathcal{P}_e$ is a clique. Therefore $L_{max}(\mathcal{P})$ is a lower bound to the size of the maximum clique of the intersection graph.

In this work we focus on the case where G is a path, i.e. the intersection graph of $\mathcal{P}$ is an *interval graph*. It is well known that every clique of an interval graph corresponds to some $\mathcal{P}_e$, therefore $L_{max}(\mathcal{P})$ is equal to the size of the maximum clique. A set of paths that no two of them intersect is an independent set of the intersection graph. When we say that a set of paths is a clique (or an independent set) we implicitly refer to their intersection graph.

Colors and Colorings: In addition to the graph G and the set $\mathcal{P}$ of paths, we are given an integer W that denotes the number of colors available. For two integers i, j such that $i \leq j$, $[i, j] \stackrel{def}{=} \{k \in \mathbb{N} : i \leq k \leq j\}$. The set of available colors is $\Lambda = [0, W-1]$. A set $[i, j] \subseteq \Lambda$ is said to be an *interval* of colors. When $0 \leq j < i \leq W-1$ we define $[i, j] \stackrel{def}{=} [i, W-1] \cup [0, j]$. In both cases $[i, j]$ is termed a *circular interval* of colors, i.e. colors that are consecutive on a ring (in which 0 is the successor of $W-1$).

A *(multi)coloring* is a function $w : \mathcal{P} \mapsto 2^\Lambda$ that assigns to each path $P \in \mathcal{P}$ a subset of the set Λ of colors. A coloring w is *valid* if for any two paths $P, P' \in \mathcal{P}$ such that $P \cap P' \neq \emptyset$ we have $w(P) \cap w(P') = \emptyset$. For a color $\lambda \in \Lambda$, $\mathcal{P}^w_\lambda$ denotes the set of paths assigned the color λ by w, i.e. $\mathcal{P}^w_\lambda = \{P \in \mathcal{P} : \lambda \in w(P)\}$. If w is a valid coloring, then for any two paths $P, P' \in \mathcal{P}^w_\lambda$ we have $P \cap P' = \emptyset$. In other words, $\mathcal{P}^w_\lambda$ is an independent set of $\mathcal{P}$. When there is no ambiguity, we omit the superscript w and denote $\mathcal{P}^w_\lambda$ as $\mathcal{P}_\lambda$.

A coloring is *contiguous* (resp. *circularly contiguous*), if for every $P \in \mathcal{P}$, $w(P)$ is an interval (resp. circular interval) of colors.

Vector notation and profits: Throughout the paper we use *vectors* of integers indexed by the elements of $\mathcal{P}$. We denote vectors with bold typeface, e.g. $\mathbf{v} = \{\mathbf{v}_P : P \in \mathcal{P}\}$. The vector $\mathbf{0}$ is the zero vector, $\mathbf{1}$ is the vector consisting of a 1 in every index.

The size vector of a coloring w is a vector $\mathbf{s}(w)$ such that $\mathbf{s}(w)_P = |w(P)|$ for every $P \in \mathcal{P}$, i.e. the entries of $\mathbf{s}(w)$ are the number of colors assigned to each path. We say that a coloring w is a $(\mathbf{a}-\mathbf{b})$-*coloring* if $\mathbf{a} \leq \mathbf{s}(w) \leq \mathbf{b}$, and w is a $\mathbf{v}$-*coloring* if it is a $(\mathbf{v}-\mathbf{v})$-*coloring*. An ordinary coloring in which every path is assigned one color corresponds to a $\mathbf{1}$-*coloring*, and clearly any coloring is a $(\mathbf{0} - W \cdot \mathbf{1})-$coloring.

Given a real vector $\mathbf{w}$ of *weights*, the *profit* $p^w(P, \mathbf{w})$ obtained by a coloring w, from a path P is $p^w(P, \mathbf{w}) \stackrel{def}{=} \mathbf{w}_P \cdot |w(P)|$. The total profit due to a coloring w is $p^w(\mathcal{P}, \mathbf{w}) \stackrel{def}{=} \sum_{P \in \mathcal{P}} p^w(P, \mathbf{w})$.

In this work we use the term *maximum independent set* to mean an independent set with maximum profit, and denote the profit obtained from such a set as $\alpha(\mathcal{P}, \mathbf{w})$. Usually the weight function under consideration will be clear from the context and we will use $p^w(\mathcal{P})$ (resp. $\alpha(\mathcal{P})$) as a shorthand for $p^w(\mathcal{P}, \mathbf{w})$ (resp. $\alpha(\mathcal{P}, \mathbf{w})$).

We note that $p^w(\mathcal{P}) = \sum_{P \in \mathcal{P}} p^w(P) = \sum_{P \in \mathcal{P}} \mathbf{w}_P \cdot |w(P)| = \mathbf{w} \cdot \mathbf{s}(w)$. We can write the profit of a valid coloring w, from a path P as the sum of the profits obtained from every color of P, i.e. $p^w(P) = \mathbf{w}_P \cdot |w(P)| = \sum_{\lambda \in w(P)} \mathbf{w}_P$ and therefore

$$p^w(\mathcal{P}) = \sum_{P \in \mathcal{P}} \sum_{\lambda \in w(P)} \mathbf{w}_P = \sum_{\lambda \in \Lambda} \sum_{P \in \mathcal{P}^w_\lambda} \mathbf{w}_P \leq \sum_{\lambda \in \Lambda} \alpha(\mathcal{P}) = W \cdot \alpha(\mathcal{P})$$

where the inequality follows from the fact that $\mathcal{P}^w_\lambda$ is an independent set.

The Problem(s): In this work we consider the following problem and its variants.

> PROFIT MAXIMIZING COLORING(PMC)
> `Input:` A tuple $(G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w})$ where G is a graph, $\mathcal{P}$ is a set of paths on G, W is an integer, $\mathbf{a}$ and $\mathbf{b}$ are two integer vectors and $\mathbf{w}$ is a real vector indexed by $\mathcal{P}$.
> `Output:` A valid $(\mathbf{a}-\mathbf{b})$-*coloring* w.
> `Objective:` Maximize $p^w(\mathcal{P}, \mathbf{w})$.

The problems Profit Maximizing Contiguous Coloring (PMCC) and Profit Maximizing Circularly Contiguous Coloring (PMCCC) problems are variants of PMC in which the coloring w has to be contiguous, and circularly contiguous, respectively.

We denote the optimum of an instance $(G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w})$ of a problem PRB $\in$ {PMC, PMCC, PMCCC} by $\text{OPT}_{\text{PRB}}(G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w})$. A contiguous coloring is a circularly contiguous coloring, which is in turn a coloring. Therefore we have:

$$\begin{aligned}\textsc{Opt}_{\textsc{Pmcc}}(G,\mathcal{P},W,\mathbf{a},\mathbf{b},\mathbf{w}) &\leq \textsc{Opt}_{\textsc{Pmccc}}(G,\mathcal{P},W,\mathbf{a},\mathbf{b},\mathbf{w}) \\ &\leq \textsc{Opt}_{\textsc{Pmc}}(G,\mathcal{P},W,\mathbf{a},\mathbf{b},\mathbf{w}).\end{aligned} \tag{1}$$

Any coloring, and in particular an optimal one that we denote by w^*, satisfies $p^{w^*}(\mathcal{P}) \leq W \cdot \alpha(\mathcal{P})$. Therefore we have

$$\textsc{Opt}_{\textsc{Pmc}}(G,\mathcal{P},W,\mathbf{a},\mathbf{b},\mathbf{w}) \leq W \cdot \alpha(\mathcal{P}).$$

We now observe that the above inequalities are tight when the lower and upper bounds $\mathbf{a}$ and $\mathbf{b}$ are trivial. In other words, in this case all the three problems equivalent to the problem of finding $\alpha(\mathcal{P})$.

Proposition 1. *If* $\mathbf{a} = \mathbf{0}, \mathbf{b} = W \cdot \mathbf{1}$ *then*

$$\begin{aligned}\textsc{Opt}_{\textsc{Pmcc}}(G,\mathcal{P},W,\mathbf{a},\mathbf{b},\mathbf{w}) &= \textsc{Opt}_{\textsc{Pmccc}}(G,\mathcal{P},W,\mathbf{a},\mathbf{b},\mathbf{w}) \\ &= \textsc{Opt}_{\textsc{Pmc}}(G,\mathcal{P},W,\mathbf{a},\mathbf{b},\mathbf{w}) = W \cdot \alpha(\mathcal{P}).\end{aligned}$$

Proof. It suffices to show that $\textsc{Opt}_{\textsc{Pmcc}}(G,\mathcal{P},W,\mathbf{a},\mathbf{b},\mathbf{w}) \geq W \cdot \alpha(\mathcal{P})$. Indeed, let $\mathcal{I}$ be a maximum independent set of $\mathcal{P}$. The coloring that assigns Λ to ever path of $\mathcal{I}$ and $\emptyset$ to all the rest is a valid contiguous $(\mathbf{0} - W \cdot \mathbf{1})$-coloring with profit $W \cdot \alpha(\mathcal{P})$. □

Path Networks: When G is a path we assume without loss of generality that the vertex set of G is $[1, n]$ where the vertices are numbered according to their order in G. We sometimes refer to the vertices and edges of G as drawn on the real line where 1 is the leftmost vertex and n is the rightmost one. Given this numbering, $s(P)$ and $t(P)$ denote the endpoints of a path P with $s(P) < t(P)$. We term these vertices as the *start* and *termination* vertices of P, respectively. We denote a sub-path of G with endpoints $i < j$ as $[i, j]$, i.e. $P = [s(P), t(P)]$. Given a sub-path δ of G, $\mathcal{P}_\delta$ denotes the set of all paths of $\mathcal{P}$ that are contained in δ.

3 Profit Maximizing Colorings

A maximum independent set can be calculated in polynomial time when the network is a path [6]. By Proposition 1 this implies an algorithm for all three problems for the case where G is a path and $\mathbf{a} = \mathbf{0}$ and $\mathbf{b} = W \cdot \mathbf{1}$. In this section we extend the study to path networks for arbitrary $\mathbf{a}$ and $\mathbf{b}$, and provide a polynomial-time optimal algorithm.

We first introduce notations and definitions that we use in this section. Let w be a coloring of a set $\mathcal{Q}$ of paths, and $\mathcal{Q}' \subseteq \mathcal{Q}$. $w' = w\big|_{\mathcal{Q}'}$ denotes the coloring w restricted to $\mathcal{Q}'$, i.e. $w'(P) = w(P)$ whenever $P \in \mathcal{Q}'$, and $w'(P) = \emptyset$ otherwise.

We reduce $\textsc{Pmc}$ to the Minimum Cost Maximum Flow ($\textsc{MinCostMaxFlow}$) problem that is well known to be solvable in polynomial time [10]. Instances of

MinCostMaxFlow are tuples $(H, s, t, \kappa, \kappa', c)$ where H is a directed graph, $s \in V(H)$ (resp. $t \in V(H)$) is the source (resp. sink) vertex, $\kappa : A(H) \mapsto \mathbb{R}$ (resp. $\kappa' : A(H) \mapsto \mathbb{R}$) determines the lower (resp. upper) bounds of the flow on every arc, and finally $c : A(H) \mapsto \mathbb{R}$ determines the cost of a unit flow on every arc. The goal is to find a flow $f : A(H) \mapsto \mathbb{R}$ from s to t that has a minimum cost among all maximum flows, i.e. among all flows of maximum amount, as follows. Recall that the amount of a flow f is the amount of flow entering t, i.e. $\sum_{e \in \delta_H^-(t)} f(e)$ and its cost $c(f)$ is $\sum_{e \in A(H)} f(e) \cdot c(e)$.

Given an instance $I = (G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w})$ of Pmc, we build a flow network $N(I) = (H, s, t, \kappa, \kappa', c)$. For convenience we introduce two additional semi-infinite (i.e. having one endpoint) paths $P^{(-)} = [-\infty, 1]$ and $P^{(+)} = [n, \infty]$ with zero profit, and we define $\mathcal{P}' = \mathcal{P} \cup \{P^{(-)}, P^{(+)}\}$. $V(H) = S \cup T$ where $T = \{t_P : P \in \mathcal{P}'\}, S = \{s_P : P \in \mathcal{P}'\}$. $A(H) = A_1 \cup A_2$ where $A_1 = \{(s_P, t_P) : P \in \mathcal{P}'\}$ and $A_2 = \{(t_P, s_{P'}) : s(P') \geq t(P)\}$. We proceed with the bounds and costs of the arcs. For every path $P \in \mathcal{P}$ the bounds and costs on the corresponding arc $a = (s_P, t_P) \in A_1$ are $\kappa(a) = \mathbf{a}_P$, $\kappa'(a) = \mathbf{b}_P$ and $c(a) = -\mathbf{w}_P$. For each one of the two arcs a corresponding to the two semi-infinite paths we set $\kappa(a) = 0$, $\kappa'(a) = W$ and $c(a) = 0$. For an arc $a = (t_P, s_{P'})$ of A_2 we set $\kappa(a) = 0$, $\kappa'(a) = \infty$ and $c(a) = 0$. Finally we set $s = s_{P^{(-)}}$ and $t = t_{P^{(+)}}$.

Lemma 1. *For every feasible coloring w of an instance I of* Pmc*, there is a maximum flow $f^{(w)}$ of $N(I)$, such that $c(f^{(w)}) = -p^w(\mathcal{P})$. Moreover, given a maximum flow f of $N(I)$ a coloring w such that $f^{(w)} = f$ can be found in polynomial-time.*

Proof. We first observe that the maximum flow of $N(I)$ is W. Indeed a flow of amount W can be pushed from $s_{P^{(-)}}$ via $t_{P^{(-)}}$ and $s_{P^{(+)}}$ to $t_{P^{(+)}}$. On the other hand this is a maximum flow because the arc $(s_{P^{(-)}}, t_{P^{(-)}})$ constitutes a cut of weight W.

Given a feasible coloring w of I we define the flow $f^{(w)}$ as the sum of W flows $f_1^{(w)}, f_2^{(w)}, \ldots, f_W^{(w)}$. For each color $\lambda \in \Lambda$, $f_\lambda^{(w)}$ corresponds to the independent set $\mathcal{P}_\lambda^w$. $f_\lambda^{(w)}$ pushes one unit of flow from $s_{P^{(-)}}$ to $t_{P^{(+)}}$ over the path that consists of the arcs of A_1 corresponding to the paths of $\mathcal{P}_\lambda^w$ and the arcs of A_2 connecting two consecutive paths of $\mathcal{P}_\lambda^w$. The cost of an A_2 arc is zero, and the cost of an A_1 arcs corresponding to a path P is $-\mathbf{w}_P$. Therefore the cost of f_λ is $c(f_\lambda^{(w)}) = -\sum_{P \in \mathcal{P}_\lambda^w} \mathbf{w}_P$. Summing up over all colors λ we get

$$c(f^{(w)}) = \sum_{\lambda \in \Lambda} c(f_\lambda^{(w)}) = -\sum_{\lambda \in \Lambda} \sum_{P \in \mathcal{P}_\lambda^w} \mathbf{w}_P = -p^w(\mathcal{P}).$$

$f^{(w)}$ satisfies the bounds κ and κ'. Indeed, for an arc a of A_1 corresponding to a path $P \in \mathcal{P}'$ we have $f^{(w)}(a) = |w(P)|$ and $\kappa(a) = \mathbf{a}_P \leq |w(P)| \leq \mathbf{b}_P = \kappa'(a)$. For the arcs of A_2 we have $\kappa(a) = 0 \leq f^{(w)}(a) \leq \infty = \kappa'(a)$.

Any maximum flow f of $N(I)$ can be split, in polynomial time, into W unit flows $f_1, f_2, \ldots, f_W$. Each unit flow uses a path from $s_{P^{(-)}}$ to $t_{P^{(+)}}$. Such a path starts with an A_1 arc, and alternates between A_1 and A_2 arcs. The set of odd arcs corresponds to an independent set paths of $\mathcal{P}$. w is defined such that $\mathcal{P}^w_\lambda$ is the independent set corresponding to f_λ. $\square$

Corollary 1. *The profit $p^w(\mathcal{P})$ is maximized when $c(f^{(w)})$ is minimum.*

This implies the following a polynomial time algorithm for Pmc: Given an instance I, calculate a minimum cost maximum flow f of $N(I)$ and return a coloring w such that $f^{(w)} = f$.

4 Profit Maximizing Contiguous Colorings

In this section we consider contiguous colorings. We first observe that the problem is Np-Hard even if the graph is a path. In Section 4.1 we compare circularly contiguous colorings to contiguous colorings and we provide an algorithm that transforms a circularly contiguous coloring to a contiguous coloring with a small loss in the profit. In Section 4.2 we consider the case where the load on the edges is bounded by some constant and provide a polynomial-time algorithm for this case. In Section 4.3 we provide an approximation algorithm for another special case where the paths constitute a proper set.

Let G be a graph and f a weight function $f : V(G) \to \mathbb{N}$ on its vertices. An interval coloring w of G, f assigns an interval $w(v)$ of $f(v)$ integers to every vertex v of G, such that $f(v) \cap f(v') = \emptyset$ whenever v and v' are adjacent in G. The weight $f(K)$ of a clique $K \subseteq V(G)$ is the sum $\sum_{v \in K} f(v)$ of the individual weights of its vertices. The clique number $\omega(G, f)$ of the weighted graph (G, f) is the maximum weight of its cliques. The interval chromatic number of $\chi(G, f)$ is the minimum number of colors used by an interval coloring of (G, f) [6]. Clearly $\chi(G, f) \geq \omega(G, f)$. The problem of finding the interval chromatic number of a weighted interval graph is also known as the *shipbuilding* problem, and also as the *dynamic storage allocation* problem. This problem is known to be Np-Complete [7]. Therefore

Lemma 2. Pmcc *is* Np-*Hard even when G is a path.*

Proof. Let (G, f) be a weighted interval graph, and $\mathcal{P}$ the set of paths on a path H which represent G. Let $\mathbf{w}$ be any weight function on $\mathcal{P}$. The instance $(H, \mathcal{P}, W, f, f, \mathbf{w})$ is feasible if and only if the interval chromatic number of (G, f) is at most W. $\square$

4.1 Comparison with Circularly Contiguous Colorings

In this section we present the algorithm CircularToContiguous that converts a circularly contiguous $(\mathbf{a}-\mathbf{b})$-*coloring* w^{cc} to a contiguous $(\lceil \mathbf{a/2} \rceil - \mathbf{b})$-*coloring* w^c such that $p^{w^c}(\mathcal{P}) \geq \frac{3}{4} p^{w^{cc}}(\mathcal{P})$.

A circularly contiguous interval $[i, j]$ is either contiguous or the disjoint union of two contiguous intervals $[j, W-1]$, $[0, i]$. The size of one of these sub-intervals is at least half of the size of the entire interval. CircularToContiguous chooses a color $\bar{\lambda}$ uniformly at random and renames all the colors such that $\bar{\lambda}$ becomes 0, $(\bar{\lambda}+1) \mod W$ becomes 1, and so on. Then to every path P for which the obtained coloring is not contiguous it assigns the biggest among the two corresponding contiguous colorings.

w^c is clearly a contiguous $(\lceil \mathbf{a/2} \rceil - \mathbf{b})$-*coloring*. For a given path P we now calculate the expected value of $|w^c(P)|$. Let $\ell = |w^{cc}(P)|$, and $[i,j] = w^{cc}(P)$. We consider three cases: (a) $\bar{\lambda}$ is not in $[i+1,j]$. In this case, after the renaming phase, $w^c(P)$ is contiguous. Therefore $|w^c(P)| = \ell$. (b) $\bar{\lambda} = i+k$ and $k <= \ell/2$. In this case $|w^c(P)| = \ell - k$. (c) $\bar{\lambda} = i+k$ and $l/2 < k < l$. In this case $|w^c(P)| = k$. The probability that $\bar{\lambda}$ gets any given value is 1/W. We consider only the case that ℓ is even which leads to a smaller expected value. We have

$$
\begin{aligned}
E[|w^c(P)|] &= \frac{1}{W}\left(\sum_{k=1}^{\ell/2}(\ell-k) + \sum_{k=\ell/2+1}^{l-1} k + (W-\ell+1)\ell\right) \\
&= \frac{1}{W}\left(\frac{3}{4}\ell^2 - \ell + (W-\ell+1)\ell\right) = \ell - \frac{\ell}{W}\frac{\ell}{4} \geq \frac{3}{4}\ell = \frac{3}{4}\left|w^{cc}(P)\right|.
\end{aligned}
$$

We use the above inequality and linearity of expectation to calculate the expected value of the solution.

$$
E[p^{w^c}(\mathcal{P})] = E[\mathbf{w}\cdot\mathbf{s}(w^c)] = \mathbf{w}\cdot E[\mathbf{s}(w^c)] \geq \frac{3}{4}E[\mathbf{s}(w^{cc})] = \frac{3}{4}p^{w^{cc}}(\mathcal{P}).
$$

Therefore

Lemma 3. *There is a randomized polynomial-time algorithm that converts a valid circularly contiguous* $(\mathbf{a}-\mathbf{b})$*-coloring* w^{cc} *to a valid contiguous* $(\lceil \mathbf{a/2} \rceil - \mathbf{b})$*-coloring* w^c *satisfying* $p^{w^c}(\mathcal{P}) \geq \frac{3}{4}p^{w^{cc}}(\mathcal{P})$.

The above randomized algorithm can be de-randomized by trying every possible value of W and picking up the best result. Clearly at least one solution is at least as good as the expected value. This de-randomization does not lead to a polynomial-time algorithm whenever the value of W is exponential in the input size. An efficient de-randomization can be obtained by guessing each one bit of $\bar{\lambda}$ at a time. We conclude

Lemma 4. *There is a deterministic polynomial-time algorithm that converts a valid circularly contiguous* $(\mathbf{a}-\mathbf{b})$*-coloring* w^{cc} *to a valid contiguous* $(\lceil \mathbf{a/2} \rceil - \mathbf{b})$*-coloring* w^c *satisfying* $p^{w^c}(\mathcal{P}) \geq \frac{3}{4}p^{w^{cc}}(\mathcal{P})$.

4.2 Bounded Load

Let $I = (G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w})$ be an instance of Prb $\in$ {Pmc, Pmcc, Pmccc}, and let $v \in [1,n]$. We denote by $I^{(v+)}$ the instance obtained from I by restricting the paths set to ones that start at vertex v or before. Formally $I^{(v+)} =$

$(G, \mathcal{P}^{(v+)}, W, \mathbf{a}^{(v)}, \mathbf{b}^{(v+)}, \mathbf{w}^{(v+)})$ where $\mathcal{P}^{(v+)} = \{P \in \mathcal{P} : s(P) \leq v\}$, $\mathbf{a}^{(v)} = \mathbf{a}\big|_{\mathcal{P}^{(v+)}}$, $\mathbf{b}^{(v+)} = \mathbf{b}\big|_{\mathcal{P}^{(v+)}}$ and $\mathbf{w}^{(v+)} = \mathbf{w}\big|_{\mathcal{P}^{(v+)}}$.

We say that two colorings w, w' of two subsets $\mathcal{Q}, \mathcal{Q}'$ of $\mathcal{P}$ *agree* if $w(P) = w'(P)$ whenever $P \in \mathcal{Q} \cap \mathcal{Q}'$, and we denote this by $w \sim w'$. Let $\bar{w}$ be a coloring of the paths $\mathcal{P}_{e_v}$ where e_v denotes the edge $\{v-1, v\}$. We denote by $\textsc{Opt}_{\textsc{Prb}}(I, v, \bar{w})$ the optimum of problem $\textsc{Prb}$ for the instance $I^{(v)}$ when the feasible colorings are restricted to colorings that *agree* with $\bar{w}$. As our goal in this section is to provide an optimal algorithm for $\textsc{Pmccc}$, and this implies an optimal algorithm for all problems, in the sequel we refer only to this problem, although the arguments hold for all three problems. Clearly

$$\textsc{Opt}_{\textsc{Pmcc}}(I) = \textsc{Opt}_{\textsc{Pmcc}}(I^{(n)}) = \max\left\{\textsc{Opt}_{\textsc{Pmc}}(I, n, \bar{w}) : \bar{w} \text{ is a cont. coloring of } \mathcal{P}_{e_n}\right\}.$$

Consider a contiguous coloring w of $\mathcal{P}^{(v+)}$, and a contiguous coloring w^- of $\mathcal{P}^{(v-1)}$ that agrees w. We have

$$p^w(\mathcal{P}^{(v+)}) = p^{w^-}(\mathcal{P}^{(v-1)}) + \sum_{P \text{ s.t. } s(P)=v-1} |w(P)| \cdot \mathbf{w}_P.$$

We note that the second term depends only on $w\big|_{\mathcal{P}_{e_v}}$. Among all contiguous colorings w that agree with a given contiguous coloring $\bar{w}$ of $\mathcal{P}_{e_v}$, the second term is a constant. Therefore the maximum is obtained at the maximum of the first term. We conclude

$$\textsc{Opt}_{\textsc{Pmcc}}(I, v, \bar{w}) = \max_{\bar{w}^- \sim \bar{w}} \left\{\textsc{Opt}_{\textsc{Pmcc}}(I, v-1, \bar{w}^-)\right\} + \sum_{P \text{ s.t. } s(P)=v-1} |w(P)| \cdot \mathbf{w}_P.$$

These equations imply the dynamic programming algorithm $\textsc{ContColorDynProg}$. For simplicity $\textsc{ContColorDynProg}$ calculates the optimum of the instance without explicitly finding an optimal coloring. It can be easily extended to return an optimal coloring.

The loops at lines 3 and 6 constitute the dominant part in the running time of the algorithm. A contiguous coloring of $\mathcal{P}_{e_v}$ can be found by fixing a permutation of the $\ell = L_{e_v}$ paths, and assigning to each path a positive number so that their sum does not exceed W. The number of permutations is $\ell!$ and the number of possible assignments of the numbers is $\binom{W}{\ell}$. Therefore each one of the loops iterates at most $\ell!\binom{W}{\ell} \leq W^\ell$ times, and the total number of iterations is at most $W^{2\cdot\ell} \leq W^{2\cdot L_{max}(\mathcal{P})}$. Therefore

Lemma 5. *There is a polynomial-time algorithm that solves* $\textsc{Pmc}(G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w})$ *when G is a path network and $L_{max}(\mathcal{P})$ is bounded by a constant.*

4.3 Proper Sets of Paths

A set of paths is *proper* if no path in the set properly contains another. The intersection graph of a proper set of paths on a path graph is a *proper interval*

Algorithm 1 $\textsc{ContColorDynProg}\ I = (G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w})$

$\textsc{Opt}_{\textsc{Pmc}}(I, 1, w_{empty}) \leftarrow 0.$ ▷ w_{empty} is the empty coloring.
for $v = 2$ to $v = n$ **do**
  for all Contiguous colorings $\bar{w}$ of $\mathcal{P}_{e_v}$ **do**
    $C \leftarrow \sum_{P \text{ s.t. } s(P)=v-1} |w(P)| \cdot \mathbf{w}_P.$
    $M \leftarrow 0.$
    for all Contiguous colorings $\bar{w}^-$ of $\mathcal{P}_{e_{v-1}}$ s.t. $\bar{w}^- \sim \bar{w}$ **do**
      if $\textsc{Opt}_{\textsc{Pmcc}}(I, v-1, \bar{w}^-) > M$ **then**
        $M \leftarrow \textsc{Opt}_{\textsc{Pmcc}}(I, v-1, \bar{w}^-).$
      end if
    end for
    $\textsc{Opt}_{\textsc{Pmcc}}(I, v, \bar{w}) \leftarrow M + C.$
  end for
end for
return $\max\{\textsc{Opt}_{\textsc{Pmc}}(I, n, \bar{w}) : \bar{w}$ is a contiguous coloring of $\mathcal{P}_{e_n}\}$.

graph. Let P, P' be two paths in a proper set $\mathcal{P}$ of paths. $s(P) \leq s(P')$ if and only of $t(P) \leq t(P')$.

We present a simple algorithm $\textsc{ProperToCircular}$ that converts any coloring w of a proper set of paths can to a circularly contiguous coloring w^{cc} with the same profit. $\textsc{ProperToCircular}$ iterates over the paths according to the total order implied by their start vertices. Every path is assigned a circular interval $[\lambda, \lambda + |w(P)| - 1]$ where $\lambda = 0$ for the first path, and for each subsequent path λ is the last color of the previous path, plus one. Clearly w^{cc} is a circularly contiguous $(\mathbf{a}-\mathbf{b})$-*coloring* and $p^{w^{cc}}(\mathcal{P}) = p^w(\mathcal{P})$. It remains to show that w^{cc} is valid.

Assume, by way of contradiction, that w^{cc} is not valid. Then there are two intersecting paths $P, P' \in \mathcal{P}$ and a color λ such that $\lambda \in w^{cc}(P) \cap w^{cc}(P')$. Assume without loss of generality that $s(P) \leq s(P')$, and let e be the last edge of P, i.e. $e = \{t(P)-1, t(P)\}$. As $\mathcal{P}$ is a proper set of paths and $P \cap P' \neq \emptyset$, we have $e \in P'$. Moreover any path P'' such that $s(P) \leq s(P'') \leq s(P')$ contains the edge e. Therefore the set $\mathcal{Q}$ of all paths whose start vertices are between $s(P)$ and $s(P')$ (inclusive) is a subset of $\mathcal{P}_e$. As $\textsc{ProperToCircular}$ considers the paths in the order of their start vertices, and λ was used in both P and P', this means that the number of colors assigned by w^{cc} to the paths of $\mathcal{Q}$ exceeds W. However, this is exactly the number of colors assigned to these paths by w. Then w assigns more than W colors to the paths of $\mathcal{P}_e$, therefore invalid, contradicting our assumption.

Combining with (1) we conclude

Lemma 6. *When G is a path and $\mathcal{P}$ is a proper set of paths*

$$\textsc{Opt}_{\textsc{Pmccc}}(G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w}) = \textsc{Opt}_{\textsc{Pmc}}(G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w}).$$

Moreover there is a polynomial-time algorithm solving $\textsc{Pmccc}(G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w})$ *optimally.*

Combining this with Lemma 4 we obtain the following two corollaries.

Corollary 2. *There is a deterministic polynomial-time* $4/3$*-approximation algorithm for* PMCC $(G, \mathcal{P}, W, \mathbf{a}, \mathbf{b}, \mathbf{w})$ *when* G *is a path,* $\mathcal{P}$ *is a proper set of paths,* $\mathbf{b}$ *is a valid coloring and* $\mathbf{a} \leq \lceil \mathbf{b}/2 \rceil$.

5 Future Work

A few open problems regarding contiguous colorings in path networks that are closely related to our results, namely: a) to find an approximation algorithm for PMCC, b) to obtain prove APX-hardness of PMCC, c) To determine if PMCC is polynomial time solvable for proper intervals.

Another research direction is to extend the results to other topologies, especially those that are relevant in optical networks, such as rings, trees, grids, bounded treewidth. Finally, as stated in the introduction, the flex-grid technology opens a wide range of problems, such as regenerator placement, traffic grooming etc., that have been studied in the fixed-grid context, to be reconsidered in the flex-grid context.

References

1. R. Ramaswami, K. N. Sivarajan, and G. H. Sasaki. *Optical Networks: A Practical Perspective*. Morgan Kaufmann Publisher Inc., San Francisco, 2009.
2. R. Klasing. Methods and problems of wavelength-routing in all-optical networks. In *Proceeding of the MFCS'98 Workshop on Communication, August 24-25, Brno, Czech Republic*, pages 1–9, 1998.
3. M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, and S. Matsuoka. Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies. *Comm. Mag.*, 47:66–73, nov. 2009.
4. O. Gerstel. Realistic approaches to scaling the IP network using optics. In *Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC)*, pages 1 –3, march 2011.
5. O. Gerstel. Flexible use of spectrum and photonic grooming. In *Photonics in Switching, OSA (Optical Society of America) Technical Digest*, page paper PMD3, 2010.
6. Martin Charles Golumbic. *Algorithmic Graph Theory and Perfect Graphs (Annals of Discrete Mathematics, Vol 57)*. North-Holland Publishing Co., Amsterdam, The Netherlands, The Netherlands, 2004.
7. M. Garey and D. S. Johnson. *Computers and Intractability, A Guide to the Theory of NP-Completeness*. Freeman, 1979.
8. Magnús M. Halldórson and Guy Kortzarz. Multicoloring: Problems and techniques. In *29th International Symposium on Mathematical Foundations of Computer Science, MFCS*, LNCS 3153, pages 25–41. Springer-Verlag, August 2004.
9. Adam L. Buchsbaum, Howard Karloff, Claire Kenyon, Nick Reingold, and Mikkel Thorup. Opt versus load in dynamic storage allocation. *SIAM Journal of Computing*, 33(3):632–646, 2004.
10. Andrew V. Goldberg and Robert E. Tarjan. Finding minimum-cost circulations by canceling negative cycles. *J. ACM*, 36(4):873–886, October 1989.