

On a Non-cooperative Model for Wavelength Assignment in Multifiber Optical Networks^{*}

Evangelos Bampas, Aris Pagourtzis, George Pierrakos, and Katerina Potika

School of Elec. & Comp. Eng., National Technical University of Athens
Polytechniupoli Zografou, 157 80 Athens, Greece
{ebamp, pagour, mpierr, epotik}@cs.ntua.gr

Abstract. We study path multicoloring games that describe situations in which selfish entities possess communication requests in a multifiber all-optical network. Each player is charged according to the maximum fiber multiplicity that her color (wavelength) choice incurs and the social cost is the maximum player cost. We investigate the price of anarchy of such games and provide two different upper bounds for general graphs—namely the number of wavelengths and the minimum length of a path of maximum disutility, over all worst-case Nash Equilibria—as well as matching lower bounds which hold even for trees; as a corollary we obtain that the price of anarchy in stars is exactly 2. We also prove constant bounds for the price of anarchy in chains and rings in which the number of wavelengths is relatively small compared to the load of the network; in the opposite case we show that the price of anarchy is unbounded.

Key words: selfish wavelength assignment, non-cooperative games, price of anarchy, multifiber optical networks, path multicoloring

1 Introduction

The need for efficient access to the optical bandwidth in all-optical networks has given rise to the study of several optimization problems in the past years. The most well-studied among them is the problem of assigning a path and a color (wavelength) to each communication request in such a way that paths of the same color are edge-disjoint and the number of colors used is minimized. Nonetheless, it has become clear that the number of wavelengths in commercially available fibers is rather limited—and will probably remain such in the foreseeable future. Fortunately, the use of multiple fibers has come to the rescue. However, fibers are not unlimited either, therefore it makes sense to minimize their usage. This is particularly interesting from the customer's point of view, for example in

^{*} This work has been funded by the project PENED 2003. The project is cofinanced 75% of public expenditure through EC–European Social Fund, 25% of public expenditure through Ministry of Development–General Secretariat of Research and Technology of Greece and through private sector, under measure 8.3 of Operational Programme “Competitiveness” in the 3rd Community Support Programme.

situations where one can hire a number of parallel fibers for a certain period and the cost depends on that number.

To this end, several optimization problems have been defined and studied, the objective being to minimize either the maximum fiber multiplicity per edge [1–3] or the sum of these maximum multiplicities over all edges of the graph [4–6]; in another scenario the allowed fiber multiplicity per edge is given and the goal is to minimize the number of wavelengths needed [7, 8, 5].

In this work we consider a non-cooperative model, where each request is issued by a user who tries to optimize her own fiber usage by selecting the most appropriate wavelength, taking into account other users’ choice. This model is mainly motivated by the lack of centralized control in large scale networks. We assume that each user is charged according to the maximum fiber multiplicity that the user’s choice incurs. More specifically, a user will be charged according to the maximum number of paths that share an edge with her and use the same wavelength. We consider as *social cost* the maximum fiber multiplicity that appears on any edge of the network. Minimizing this quantity is particularly important in cases where fibers are hired or sold as a whole, hence the maximum number of fibers needed on an edge determines the total cost; further motivation can be found in papers that address the corresponding optimization problem (see e.g. [1–3]). Here we focus on situations where routing is unique (acyclic topologies) or pre-determined—as happens in many practical settings, for example in cases where there are specific routing constraints such as a requirement to use lightpaths that have been set in advance, or shortest path routing.

We formulate the above model by defining the class of SELFISH PATH MULTICOLORING (S-PMC) games: the input is a graph, a set of paths, and the number of colors w . Each player controls a path in the graph and has to choose a color for that path from $\{\alpha_1, \dots, \alpha_w\}$. A player is charged according to the maximum multiplicity of her color along her path. We consider as *social cost* the maximum color multiplicity per edge, i.e., the maximum number of paths of same color that use an edge.

Related work. Arguably, the most important notion in the theory of non-cooperative games is the *Nash Equilibrium (NE)* [9], a stable state of the game in which no player has incentive to change strategy unilaterally. A fundamental question in this theory concerns the existence of *pure Nash Equilibria (PNE)*. For various games [10–13] it has been shown that a PNE exists and can usually be found with the use of potential functions. A standard measure of the worst-case quality of Nash Equilibria relative to optimal solutions is the *price of anarchy (PoA)* [14], which has been extensively studied for load balancing games [14, 15] and other problems such as routing and facility location [10, 16]. A second known measure related to NE is the *price of stability (PoS)*, defined in [17].

S-PMC games are closely related to a variation of congestion games [18, 19] where a player’s cost is determined by her *maximum* latency instead of the usual cost which is the *sum* of her latencies. Next, we briefly explain the relation of those models to ours.

In [18] the authors study atomic routing games on networks, where each player chooses a path to route her traffic from an origin to a destination node, with the objective of minimizing the maximum congestion on any edge of her path. They show that these games always possess at least one optimal PNE (hence the PoS is 1) and that the PoA of the game is determined by topological properties of the network; in particular they show that the PoA is upper bounded by the length of the longest path in the player strategy sets and lower bounded by the length of the longest cycle. Some of our results extend to their model, since our model mimics traffic routing in the following sense: we may consider a multigraph, where we replace each edge with w parallel edges, one for each color. Each player's strategy set then consists of w different source-destination paths, corresponding to the w available colors in the original model. A further generalization is the model of Banner and Orda [19], where they introduce the notion of bottleneck games. In this model they allow arbitrary latency functions on the edges and consider both the case of splittable and unsplittable flows. They show existence, convergence and non-uniqueness of equilibria and they prove that the PoA for these games is unbounded. Both models are more general than ours; however our model fits better into the framework of all-optical networks for which we manage to provide, among others, smaller upper bounds on the PoA compared to the ones obtained by [18, 19], as well as a better convergence rate to Nash equilibria. In [20] they study similar games and give results for restricted cases, e.g. single-commodity networks.

To the best of our knowledge selfish path multicoloring games have not been studied before. Selfish path coloring in single fiber all-optical networks have been studied in [21–24]. Bilò and Moscardelli [21] consider the convergence to Nash Equilibria of selfish routing and path coloring games. Later, Bilò et al. [22] considered different information levels of local knowledge that players may have for computing their payments in the same games and give bounds for the PoA in chains, rings and trees. The existence of Nash Equilibria and the complexity of recognizing and computing a Nash Equilibrium for selfish routing and path colorings games under several payment functions are considered by Georgakopoulos et al. [23]. In [24] upper and lower bounds of the PoA for selfish path coloring with and without routing are presented under functions that charge a player only according to her own strategy.

Our results. We first give an upper bound on the convergence rate of Nash dynamics for S-PMC games, and observe that the price of stability is always equal to 1. We also show how to efficiently compute a Nash Equilibrium of minimum social cost for S-PMC games in rooted trees, i.e. trees in which each path lies entirely on a simple path from some fixed root node to a leaf. For S-PMC games in stars, we prove that a known approximation algorithm for a related optimization problem actually gives an $\frac{1}{2}$ -approximate Nash Equilibrium.

For general graphs, we obtain two upper bounds on the PoA: the first, which is not hard to show, is equal to the number of available colors. The second, which requires more involved arguments, is equal to the length of a shortest path with

maximum disutility in any worst-case NE. For both bounds we provide matching lower bounds. In fact, we prove that these bounds hold even in trees.

Then, we move on to specific network topologies and show that for S-PMC games in stars $\text{PoA} = 2$. We also provide constant bounds on the PoA in a broad class of S-PMC games in chains and rings, namely for all games with $L = \Omega(w^2)$, where w is the number of available colors and L is the maximum load among all edges of the network. On the other hand, for any $\varepsilon > 0$ we exhibit a class of S-PMC games in chains (and rings) with $L = \Theta(w^{2-\varepsilon})$ for which the PoA is unbounded.

In order to show our upper bounds, we demonstrate path patterns that must be present in any Nash Equilibrium, while for the lower bounds we employ recursive construction techniques.

2 Definitions and Model

Given an undirected graph $G(V, E)$, a set P of simple paths defined on G , and a set $W = \{\alpha_1, \dots, \alpha_w\}$ of available colors, $L(e)$ will denote the *load* of edge e , i.e. the number of paths that use edge e . The maximum of these loads will be denoted by L , i.e. $L = \max_{e \in E} L(e)$.

Given, additionally, an assignment of a color to each path we define the following:

- Definition 1.**
1. $\mu(e, c)$ will denote the multiplicity of color c on edge e , i.e. the number of paths that use edge e and are colored with color c .
 2. μ_e will denote the maximum multiplicity of any color on edge e , i.e. $\mu_e = \max_{c \in W} \mu(e, c)$.
 3. μ_{\max} will denote the maximum multiplicity of any color over all edges: $\mu_{\max} = \max_{e \in E} \mu_e$.
 4. $\mu(p, c)$ will denote the maximum multiplicity of color c over the edges of path p : $\mu(p, c) = \max_{e \in p} \mu(e, c)$.

It will be clear from the context which specific coloring we are referring to when we use the above notation.

The minimum μ_{\max} that can be attained by some coloring of the paths in P will be denoted by μ_{OPT} , i.e. $\mu_{\text{OPT}} = \min_c \mu_{\max}$ where c ranges over all possible colorings. We note immediately the following:

Fact. No coloring can achieve a μ_{\max} smaller than $\lceil \frac{L}{w} \rceil$. Thus, $\mu_{\text{OPT}} \geq \lceil \frac{L}{w} \rceil$.

We now proceed to define the class of selfish path multicoloring games and subclasses thereof.

Definition 2 (Selfish path multicoloring games). A selfish path multicoloring game is the following strategic game defined in terms of an undirected graph G , a set P of simple paths defined on G , and an integer $w > 0$:

- **Players:** there is one player for each path in P . For simplicity, we identify a player i with the corresponding path p_i .

- Strategies: a strategy for player i is a color c_i chosen from the set $W = \{\alpha_1, \dots, \alpha_w\}$ of available colors. We say that color c_i is assigned to path p_i or that path p_i is colored with color c_i . All players share the common set of strategies W .
- Disutility: given a strategy profile $\mathbf{c} = (c_1, \dots, c_{|P|})$, the disutility $f_i : W^{|P|} \rightarrow \mathbb{N}$ of each player i is defined as follows:

$$f_i(\mathbf{c}) = \mu(p_i, c_i) .$$

We denote this game by $\langle G, P, w \rangle$. The class of all selfish path multicoloring games will be denoted by S-PMC.

We will use the notation S-PMC(\mathcal{G}) to denote a subclass of S-PMC that contains only games satisfying a property \mathcal{G} (for example \mathcal{G} may constrain the graph on which the game is defined to belong to a specific graph class, etc.).

Following the standard definition, a strategy profile $\mathbf{c} = (c_1, \dots, c_{|P|})$ is said to be a *pure Nash Equilibrium* (PNE), or simply *Nash Equilibrium* (NE), if for each player i it holds that: $f_i(c_1, \dots, c'_i, \dots, c_{|P|}) \geq f_i(c_1, \dots, c_i, \dots, c_{|P|})$, for any strategy $c'_i \in W$. Moreover, following the definition of [25], we say that a strategy profile $\mathbf{c} = (c_1, \dots, c_{|P|})$ is an ε -*approximate Nash Equilibrium* if for each player i it holds that: $f_i(c_1, \dots, c'_i, \dots, c_{|P|}) \geq (1-\varepsilon) \cdot f_i(c_1, \dots, c_i, \dots, c_{|P|})$, for any strategy $c'_i \in W$.

Definition 3 (Blocking edges). If \mathbf{c} is a strategy profile for a game $\langle G, P, w \rangle$ and $p_i \in P$, we say that edge e is an α_j -blocking edge for p_i , or that it blocks α_j for p_i , if $e \in p_i$ and $\mu(e, \alpha_j) \geq f_i(\mathbf{c}) - 1$. Furthermore, the $\mu(e, \alpha_j)$ paths that are colored with α_j and use edge e are called α_j -blocking paths for p_i .

Intuitively, an α_j -blocking edge for p_i “blocks” p_i from switching to color α_j because if it did, the new disutility of path p_i would be at least $\mu(e, \alpha_j) + 1 \geq f_i(\mathbf{c})$, no better than its current choice. It is immediate from the definitions that the following property holds in any Nash Equilibrium of any S-PMC game:

Property 1 (Structural property of S-PMC Nash Equilibria). In a Nash Equilibrium, every path p must contain at least one α_j -blocking edge for p , for every color α_j .

Definition 4 (Social cost). The social cost of a strategy profile \mathbf{c} for an S-PMC game is defined as follows: $\text{sc}(\mathbf{c}) = \max_{e \in E} \mu_e = \mu_{\max}$.

It is straightforward to verify that the social cost of a strategy profile coincides with the maximum player disutility in that profile:

$$\text{sc}(\mathbf{c}) = \max_{e \in E} \mu_e = \max_{p_i \in P} f_i(\mathbf{c})$$

We define $\hat{\mu}$ to be the maximum social cost over all strategy profiles that are Nash Equilibria: $\hat{\mu} = \max_{\mathbf{c} \text{ is NE}} \text{sc}(\mathbf{c})$. Following the standard definitions, the *price of anarchy* (PoA) of a game $\langle G, P, w \rangle$ is the worst-case social cost in a Nash

Equilibrium divided by μ_{OPT} , i.e.: $\text{PoA}(\langle G, P, w \rangle) = \frac{\max_{\mathbf{c} \text{ is NE}} \text{sc}(\mathbf{c})}{\mu_{\text{OPT}}} = \frac{\hat{\mu}}{\mu_{\text{OPT}}}$. The *price of stability* (PoS) of a game is the best-case social cost in a NE divided by μ_{OPT} : $\text{PoS}(\langle G, P, w \rangle) = \frac{\min_{\mathbf{c} \text{ is NE}} \text{sc}(\mathbf{c})}{\mu_{\text{OPT}}}$. The price of anarchy (resp. stability) of a class of games $\text{S-PMC}(\mathcal{G})$ is the maximum price of anarchy (resp. stability) among all games in $\text{S-PMC}(\mathcal{G})$.

3 Existence and Computation of Nash Equilibria

We use lexicographic-order arguments similar to those in [18, 19] to show that in any S-PMC game the following holds: starting from an arbitrary strategy profile any Nash dynamics converges to a Nash Equilibrium of smaller or equal social cost. The proof is omitted.

Theorem 1. *For any game $\langle G, P, w \rangle$ in S-PMC:*

- a. *the price of stability is 1, and*
- b. *any Nash dynamics converges to a Nash Equilibrium in at most $4^{|P|}$ steps.*

Due to Theorem 1, computing a Nash Equilibrium of minimum social cost is at least as hard as the corresponding optimization problem. As noticed in [4] this problem is **NP**-hard in general graphs, in fact even in rings and stars. Therefore, it is also **NP**-hard to compute an optimal Nash Equilibrium even in the case of rings and stars. However, we show that there exists an efficient algorithm that computes optimal Nash Equilibria for a subclass of $\text{S-PMC}(\text{TREE})$. Furthermore, we show that we can use a known algorithm for $\text{PATH MULTICOLORING}$ in stars [4] to compute approximate Nash Equilibria for $\text{S-PMC}(\text{STAR})$ games. We will only state the theorems and omit the proofs.

Definition 5. *We define $\text{S-PMC}(\text{ROOTED-TREE})$ to be the subclass of S-PMC that contains games $\langle G, P, w \rangle$ with the following property: “ G is a tree and there is a node r such that each path in P lies entirely on some simple path from r to a leaf.”*

Consider the greedy algorithm that colors paths in order of non-decreasing distance from the root in such a way that the color multiplicity is the lowest possible with respect to the current partial coloring.

Theorem 2. *Given an $\text{S-PMC}(\text{ROOTED-TREE})$ game $\langle G(V, E), P, w \rangle$ with maximum load L as input, the greedy algorithm computes an optimal Nash Equilibrium of cost exactly $\lceil \frac{L}{w} \rceil$.*

Theorem 3. *There is a polynomial-time algorithm that computes a $\frac{1}{2}$ -approximate Nash Equilibrium for any $\text{S-PMC}(\text{STAR})$ game.*

4 Tight Upper Bounds for the PoA of S-PMC Games

In this section we provide two upper bounds on the PoA of any S-PMC game and we show that both of them are tight. The first bound is determined by a property of the network, namely the number of available wavelengths. The second bound is more subtle, as it depends on the length of paths with the highest disutility in worst-case Nash Equilibria. We prove that these bounds are tight even for the class S-PMC(ROOTED-TREE), and asymptotically tight for the class S-PMC(ROOTED-TREE: $\Delta = 3$), i.e. the subclass of S-PMC(ROOTED-TREE) that contains games defined on graphs with maximum degree 3.

Lemma 1. *The price of anarchy of any S-PMC game $\langle G, P, w \rangle$ is at most w .*

Proof. Let \mathbf{c} be a worst-case Nash Equilibrium of $\langle G, P, w \rangle$, hence $sc(\mathbf{c}) = \hat{\mu}$. Clearly, $\hat{\mu} \leq L$ and since the minimum social cost over all strategy profiles is $\mu_{\text{OPT}} \geq \lceil \frac{L}{w} \rceil$, it turns out that $\mu_{\text{OPT}} \geq \frac{\hat{\mu}}{w}$. This implies that $\frac{\hat{\mu}}{\mu_{\text{OPT}}} \leq w$. \square

Lemma 2. *For any worst-case Nash Equilibrium \mathbf{c} of an S-PMC game $\langle G, P, w \rangle$ and for any $p_i \in P$ with $f_i(\mathbf{c}) = sc(\mathbf{c}) = \hat{\mu}$, the price of anarchy of $\langle G, P, w \rangle$ is at most equal to the length of p_i .*

Proof. Let \tilde{e} be an edge of p_i where the color c_i chosen by p_i appears with maximum multiplicity $\hat{\mu}$: $\mu(\tilde{e}, c_i) = \hat{\mu}$. Let z denote the length of path p_i and let e_1, \dots, e_{z-1} be the edges that p uses, apart from \tilde{e} . For $1 \leq j \leq z-1$, let x_j be the number of colors that are blocked for p_i on e_j and let y be the number of colors that are blocked for p_i on \tilde{e} (since \mathbf{c} is a Nash Equilibrium, it must be that $x_1 + \dots + x_{z-1} + y \geq w-1$).

If it is the case that $z = 1$, i.e. p_i uses only edge \tilde{e} , then \tilde{e} must block all colors for p_i except c_i . This implies that the load of edge \tilde{e} is: $L(\tilde{e}) \geq \hat{\mu} + (w-1)(\hat{\mu}-1) = w\hat{\mu} - w + 1$. Therefore, the minimum social cost over all strategy profiles satisfies: $\mu_{\text{OPT}} \geq \lceil \frac{L(\tilde{e})}{w} \rceil \geq \lceil \hat{\mu} - \frac{w-1}{w} \rceil = \hat{\mu}$. We conclude that the price of anarchy in this case is equal to 1.

Now, assume that $z \geq 2$. We will prove that $L \geq 1 + \lceil \frac{w}{z} \rceil (\hat{\mu} - 1)$. First, observe that $L(\tilde{e}) \geq \hat{\mu} + y(\hat{\mu} - 1)$ and, for $1 \leq j \leq z-1$, $L(e_j) \geq 1 + x_j(\hat{\mu} - 1)$. If $y \geq \lceil \frac{w}{z} \rceil - 1$, then $L(\tilde{e}) \geq \hat{\mu} + (\lceil \frac{w}{z} \rceil - 1)(\hat{\mu} - 1) = 1 + \lceil \frac{w}{z} \rceil (\hat{\mu} - 1)$, therefore $L \geq 1 + \lceil \frac{w}{z} \rceil (\hat{\mu} - 1)$. If, on the other hand, $y < \lceil \frac{w}{z} \rceil - 1$, then $x_1 + \dots + x_{z-1} \geq w - 1 - y \geq w - \lceil \frac{w}{z} \rceil + 1$. This implies that there is some x_k such that $x_k \geq \frac{w - \lceil \frac{w}{z} \rceil + 1}{z-1} > \frac{w - \frac{w}{z} - 1 + 1}{z-1} = \frac{w}{z}$. Since x_k is an integer, it must be that $x_k \geq \lceil \frac{w}{z} \rceil$. Therefore, $L \geq L(e_k) \geq 1 + \lceil \frac{w}{z} \rceil (\hat{\mu} - 1)$.

We conclude that in any case $L \geq 1 + \lceil \frac{w}{z} \rceil (\hat{\mu} - 1)$. So, the price of anarchy is bounded as follows:

$$\text{PoA}(\langle G, P, w \rangle) = \frac{\hat{\mu}}{\mu_{\text{OPT}}} \leq \frac{\hat{\mu}}{\lceil \frac{L}{w} \rceil} \leq \frac{\hat{\mu}}{\left\lceil \frac{1 + \lceil \frac{w}{z} \rceil (\hat{\mu} - 1)}{w} \right\rceil} \leq z .$$

We omit the proof of the last inequality, which holds for all $\hat{\mu} \geq 2$, $w \geq 1$, and $z \geq 2$. \square




Fig. 1. The construction $A_z(\lambda)$ for the proof of Lemma 3. The thick lines represent the edges of the underlying graph, and the thin lines represent the paths defined on the graph. The color and multiplicity of each group of paths is written next to that group. Each shaded box represents a recursive copy of $A_z(\lambda - 1)$.

As an immediate corollary of Lemma 2, we derive the following upper bound on the price of anarchy:

Corollary 1. *The price of anarchy of any S-PMC game $\langle G, P, w \rangle$ is bounded as follows:*

$$\text{PoA} \leq \min_{e: \text{NE} \wedge \text{sc}(e) = \hat{\mu}} \min_{i: f_i(e) = \hat{\mu}} \text{length}(p_i)$$

Lemma 3. *The upper bounds of Lemma 1 and Corollary 1 are tight even for the class of S-PMC(ROOTED-TREE) games.*

Proof. We first define a recursive construction of an S-PMC game and a Nash Equilibrium for this game. The construction is illustrated in Figure 1. For any $z \geq 1$ and $\lambda \geq 1$, let $A_z(\lambda)$ be the following S-PMC game with z available colors: there are λ paths of color α_1 and length z , starting at the “root node” u_0 , which branch out into λ branches, one on each branch. Let us call these the “primary” paths for $A_z(\lambda)$. On any of the $z - 1$ edges of each such branch, one color is blocked for the primary path. The $\lambda - 1$ blocking paths of each edge branch out into an $A_z(\lambda - 1)$ game. They become primary paths for this copy of $A_z(\lambda - 1)$. The root node for the j -th recursive copy of $A_z(\lambda - 1)$ on the k -th branch is node $u_{k,j}$ (node $u_{k,1}$ is common for all branches). The base case of this recursive construction is $A_z(0)$, which is a degenerate game with no paths and no available colors, defined on a graph consisting of a single node.

Observe that for any $z \geq 1$, the construction $A_z(z)$ is an S-PMC(ROOTED-TREE) game in NE, in which all of the following are equal to z : w , L , μ_{\max} , and all path lengths. By Theorem 2, the optimal strategy profile for $A_z(z)$ has

social cost $\mu_{\text{OPT}} = \lceil \frac{L}{w} \rceil = 1$. Therefore, the ratio $\frac{\mu_{\text{max}}}{\mu_{\text{OPT}}}$ is equal to z for this Nash Equilibrium, hence the price of anarchy is at least z . \square

By appropriate modification of the construction presented in Figure 1, we obtain the following:

Lemma 4. *The upper bounds of Lemma 1 and Corollary 1 are asymptotically tight even for the class of S-PMC(ROOTED-TREE) games with maximum degree 3.*

We summarize the results of Lemmata 1, 2, 3, and 4 in the following theorem:

Theorem 4. *The price of anarchy of any S-PMC game $\langle G, P, w \rangle$ is upper-bounded both by w and by*

$$\min_{\mathbf{c}: \text{NE} \wedge \text{sc}(\mathbf{c}) = \hat{\mu}} \min_{i: f_i(\mathbf{c}) = \hat{\mu}} \text{length}(p_i) .$$

These bounds are tight for the class S-PMC(ROOTED-TREE) and asymptotically tight for the class S-PMC(ROOTED-TREE: $\Delta = 3$).

Theorem 5. *The price of anarchy of the class S-PMC(STAR) is 2.*

Proof. Lemma 2 implies an upper bound of 2 on the price of anarchy since the length of any path in a star cannot be larger than 2.

For the lower bound, it can be shown that the construction of Lemma 3 can be modified to yield a family of S-PMC(STAR) games with price of anarchy 2. More specifically every game $A_2(\lambda)$ can be embedded in a star, by using additional star rays for branching. The detailed construction is omitted. \square

5 The Price of Anarchy on Graphs of Maximum Degree 2

In this section we study the price of anarchy of path multicoloring games on chains and rings, and we prove a constant upper bound for a broad class of S-PMC(RING) games with $L = \Omega(w^2)$. Notice that this class essentially encompasses all S-PMC(RING) games of practical importance, as the number of wavelengths is limited in practice due to technological constraints, whereas L can grow large depending on network traffic. For the sake of completeness, we show that the PoA becomes quickly unbounded if we allow the network designer to provide ample wavelengths to the users, i.e. when $L = o(w^2)$.

We begin by strengthening Property 1 to prove a more involved structural property of Nash Equilibria in S-PMC(RING) games. Let $\langle G, P, w \rangle$ be an S-PMC(RING) game. Given a coloring $\mathbf{c} = (c_1, \dots, c_{|P|})$, let $P(e, \alpha_i)(\mathbf{c}) \subseteq P$ denote the set of paths colored with color α_i that use edge $e \in E$; by definition $|P(e, \alpha_i)(\mathbf{c})| = \mu(e, \alpha_i)$. For the sake of simplicity, in the rest of the section we will write $P(e, \alpha_i)$ instead of $P(e, \alpha_i)(\mathbf{c})$. Furthermore, let $[e_l, e_r]$ denote the clockwise arc starting at edge e_l and ending at edge e_r .

Lemma 5 (Structural property of S-PMC(RING) NE). *Given a game in S-PMC(RING) and a coloring \mathbf{c} thereof which is a Nash Equilibrium, for every edge e and color α_i there is an edge-simple arc $[e_l, e_r]$ with the following properties:*

- a. *for every color $\alpha_j \neq \alpha_i$, arc $[e_l, e_r]$ contains an edge which is an α_j -blocking edge for at least half of the paths in $P(e, \alpha_i)$, and*
- b. *for every edge e' of the arc $[e_l, e_r]$ it holds that $|P(e', \alpha_i) \cap P(e, \alpha_i)| \geq \left\lceil \frac{|P(e, \alpha_i)|}{2} \right\rceil$.*

Proof. Since the game is in NE, by Property 1 every path $p \in P(e, \alpha_i)$ must have at least one α_j -blocking edge, for every color $\alpha_j \neq \alpha_i$. For a fixed color $\alpha_j \neq \alpha_i$, consider the two α_j -blocking edges for some path in $P(e, \alpha_i)$ that are closest to edge e clockwise and counter-clockwise. It is not hard to see that for at least one of these two edges, call it $b(\alpha_j)$, the following property holds: the arc $[e, b(\alpha_j)]$ or the arc $[b(\alpha_j), e]$ is contained in at least $\left\lceil \frac{|P(e, \alpha_i)|}{2} \right\rceil$ of the paths in $P(e, \alpha_i)$. In case that there is only one α_j -blocking edge for all paths in $P(e, \alpha_i)$, then the property holds *a fortiori* for this edge.

For every color α_j we pick one such edge $b(\alpha_j)$. If the above property holds for arc $[e, b(\alpha_j)]$, we add $b(\alpha_j)$ to set B^+ , otherwise we add it to set B^- . We now claim that a clockwise traversal of the ring starting at edge e will first encounter all edges of B^+ and then all edges of B^- . Indeed, if one edge b^- of B^- lies before one edge b^+ of B^+ on this clockwise traversal, this would imply that b^- is traversed by the $\left\lceil \frac{|P(e, \alpha_i)|}{2} \right\rceil$ paths that contain the arc $[e, b^+]$ and thus b^- should also belong to B^+ .

The above discussion implies that if we define e_r to be the last edge of B^+ and e_l to be the first edge of B^- encountered in this clockwise traversal, then the edge-simple arc $[e_l, e_r]$ satisfies the conditions of the Lemma. \square

We now prove a constant upper bound on the price of anarchy of S-PMC(RING) games with $L = \Omega(w^2)$; denote this class by S-PMC(RING: $L = \Omega(w^2)$). This also provides an upper bound on the price of anarchy of any S-PMC(CHAIN: $L = \Omega(w^2)$) game, as every game defined on a chain can be trivially embedded in a ring topology.

We first employ the structural property of S-PMC(RING) Nash Equilibria (Lemma 5) in order to establish the existence of a heavily loaded edge in S-PMC(RING) games with $\hat{\mu} \geq w$.

Lemma 6. *In every S-PMC(RING) game $\langle G, P, w \rangle$ with $\hat{\mu} \geq w$ there is an edge with load at least $\frac{\hat{\mu}w}{4}$.*

Proof. Let $[e_l, e_r]_{P(e, \alpha_i)}$ be the arc that is obtained by applying Lemma 5 for path set $P(e, \alpha_i)$. We define P_1 to be the set of paths $P(\tilde{e}, \alpha_1)$ which induce the social cost $\hat{\mu}$. For $i \geq 2$ we define P_i to be the set of α_j -blocking paths for the path set P_{i-1} , for some color α_j not appearing at any of the path sets P_k , $k < i$, with the following property:

$$[e_l, e_r]_{P_i} \subseteq [e_l, e_r]_{P_{i-1}} \quad , \quad (1)$$

if such a path set exists. If more than one path sets with the desired property exist, we arbitrarily pick one of them.

Let e_i be the α_j -blocking edge for P_{i-1} ; based on the inductive definition of P_i as a set of blocking paths for path set P_{i-1} we can easily show that $\mu(e_i, \alpha_j) \geq \hat{\mu} - i + 1$. Applying Lemma 5(b) for color α_j and edge e_i yields the following: for every edge $e \in [e_l, e_r]_{P_i}$ we have that $\mu(e, \alpha_j) \geq \frac{\hat{\mu} - i + 1}{2}$. Furthermore, since Equation 1 holds for all $k \leq i$, the load of all edges $e \in [e_l, e_r]_{P_i}$ is at least $\sum_{\alpha_j} \mu(e, \alpha_j)$, where α_j now ranges over the colors of all path sets $P_k, k \leq i$. Hence, for every edge $e \in [e_l, e_r]_{P_i}$ we have that $L(e) \geq \sum_{\alpha_j} \mu(e, \alpha_j) \geq \sum_{k=1}^i \frac{\hat{\mu} - k + 1}{2}$.

Let now n be the first integer for which no such path set P_n exists and consider the path set P_{n-1} . Since we are at Nash Equilibrium we know that there exist α -blocking edges for paths in P_{n-1} , for every color α . We restrict our attention to the $w - n + 1$ colors, which have not yet appeared at any P_k , for $k \leq n - 1$; let α_j be one of these colors. Consider now an α_j -blocking edge e_n such that $e_n \in [e_l, e_r]_{P_{n-1}}$ (by Lemma 5(a) such an edge must exist). We now have that, at least half of the α_j -blocking paths in $P(e_n, \alpha_j)$, i.e. at least $\frac{\hat{\mu} - n + 1}{2}$ paths, extend beyond one of the edges $e_l(P_{n-1}), e_r(P_{n-1})$ of the arc $[e_l, e_r]_{P_{n-1}}$ (otherwise we would have picked $P(e_n, \alpha_j)$ to be P_n). This means that for at least half of these $w - n + 1$ blocking path sets, their paths leave the arc from the same edge, incurring on it an additional load of $\frac{w - n + 1}{2} \cdot \frac{\hat{\mu} - n + 1}{2}$.

Thus, the total load of this edge is at least $\sum_{i=1}^{n-1} \frac{\hat{\mu} - i + 1}{2} + \frac{w - n + 1}{2} \cdot \frac{\hat{\mu} - n + 1}{2} = \frac{\hat{\mu}w}{4} + (n - 1) \cdot \frac{\hat{\mu} - w + 1}{4}$. Since $\hat{\mu} \geq w$ the above sum is at least $\frac{\hat{\mu}w}{4}$. \square

Theorem 6. *The price of anarchy of any game in the class S-PMC(RING: $L = \Omega(w^2)$) is bounded by a constant.*

Proof. We distinguish between two cases:

- If $\hat{\mu} \geq w$, then by Lemma 6 we get $L \geq \frac{\hat{\mu}w}{4}$. This implies $\frac{L}{w} \geq \frac{\hat{\mu}}{4} \Rightarrow \mu_{\text{OPT}} \geq \frac{\hat{\mu}}{4} \Rightarrow \text{PoA} \leq 4$.
- If $\hat{\mu} < w$, then $\text{PoA} = \frac{\hat{\mu}}{\mu_{\text{OPT}}} \leq \frac{\hat{\mu}w}{L} < \frac{w^2}{L}$, where we used successively the facts that $\mu_{\text{OPT}} \geq \frac{L}{w}$ and $\hat{\mu} < w$. The last inequality, combined with the fact that $L = \Omega(w^2)$, implies $\text{PoA} = O(1)$. \square

Finally, we show that the price of anarchy can get arbitrarily large when the number of available colors increases; specifically, that it is unbounded for the classes S-PMC(CHAIN: $L = o(w^2)$) and S-PMC(RING: $L = o(w^2)$). The proof is omitted.

Theorem 7. *For any fixed $\varepsilon > 0$ there exists an infinite family of games in S-PMC(CHAIN: $L = \Theta(w^{2-\varepsilon})$) with $\text{PoA} = \Omega(w^{\frac{\varepsilon}{2}})$.*

References

1. Andrews, M., Zhang, L.: Minimizing maximum fiber requirement in optical networks. *J. Comput. Syst. Sci.* **72**(1) (2006) 118–131

2. Andrews, M., Zhang, L.: Complexity of wavelength assignment in optical network optimization. In: INFOCOM 2006, IEEE (2006)
3. Andrews, M., Zhang, L.: Wavelength assignment in optical networks with fixed fiber capacity. In: ICALP. LNCS 3142, Springer (2004) 134–145
4. Nomikos, C., Pagourtzis, A., Zachos, S.: Routing and path multicoloring. *Inf. Process. Lett.* **80**(5) (2001) 249–256
5. Erlebach, T., Pagourtzis, A., Potika, K., Stefanakos, S.: Resource allocation problems in multifiber WDM tree networks. In: WG. LNCS 2880, Springer (2003) 218–229
6. Winkler, P., Zhang, L.: Wavelength assignment and generalized interval graph coloring. In: SODA. (2003) 830–831
7. Margara, L., Simon, J.: Wavelength assignment problem on all-optical networks with k fibres per link. In: ICALP. LNCS 1853, Springer (2000) 768–779
8. Li, G., Simha, R.: On the wavelength assignment problem in multifiber WDM star and ring networks. *IEEE/ACM Trans. Netw.* **9**(1) (2001) 60–68
9. Nash, J.: Non-cooperative games. *The Annals of Mathematics* **54**(2) (1951) 286–295
10. Fotakis, D., Kontogiannis, S.C., Koutsoupias, E., Mavronicolas, M., Spirakis, P.G.: The structure and complexity of Nash equilibria for a selfish routing game. In: ICALP. LNCS 2380, Springer (2002) 123–134
11. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* **2** (1973) 65–67
12. Milchtaich, I.: Congestion games with player-specific payoff functions. *Games and Economic Behavior* **13** (1996) 111–124
13. Monderer, D., Shapley, L.S.: Potential games. *Games and Economic Behavior* **14** (1996) 124–143
14. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: STACS. LNCS 1563, Springer (1999) 404–413
15. Mavronicolas, M., Spirakis, P.G.: The price of selfish routing. In: STOC. (2001) 510–519
16. Roughgarden, T., Tardos, É.: How bad is selfish routing? *J. ACM* **49**(2) (2002) 236–259
17. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: FOCS, IEEE Computer Society (2004) 295–304
18. Busch, C., Magdon-Ismail, M.: Atomic routing games on maximum congestion. In: AAIM. LNCS 4041, Springer (2006) 79–91
19. Banner, R., Orda, A.: Bottleneck routing games in communication networks. In: INFOCOM 2006, IEEE (2006)
20. Caragiannis, I., Galdi, C., Kaklamanis, C.: Network load games. In: ISAAC. LNCS 3827, Springer (2005) 809–818
21. Bilò, V., Moscardelli, L.: The price of anarchy in all-optical networks. In: SIROCCO. LNCS 3104, Springer (2004) 13–22
22. Bilò, V., Flammini, M., Moscardelli, L.: On Nash equilibria in non-cooperative all-optical networks. In: STACS. LNCS 3404, Springer (2005) 448–459
23. Georgakopoulos, G.F., Kavvadias, D.J., Sioutis, L.G.: Nash equilibria in all-optical networks. In: WINE. LNCS 3828, Springer (2005) 1033–1045
24. Milis, I., Pagourtzis, A., Potika, K.: Selfish routing and path coloring in all-optical networks. In: CAAN. LNCS 4852, Springer (2007) 71–84
25. Chien, S., Sinclair, A.: Convergence to approximate Nash equilibria in congestion games. In: SODA, SIAM (2007) 169–178