On a Non-Cooperative Model for Wavelength Assignment in Multifiber Optical Networks

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

***Abstract*—We propose and investigate SELFISH PATH MULTICOLORING games as a natural model for non-cooperative wavelength assignment in multifiber optical networks. In this setting, we view the wavelength assignment process as a strategic game in which each communication request selfishly chooses a wavelength, in an effort to minimize the maximum congestion that it encounters on the chosen wavelength. We measure the cost of a certain wavelength assignment as the maximum, among all physical links, number of parallel fibers employed by this assignment.**

We start by settling questions related to the existence and computation of, and convergence to pure Nash equilibria in these games. Our main contribution is a thorough analysis of the *price of anarchy* of such games, that is the worst-case ratio between the cost of a Nash equilibrium and the optimal cost. We first provide upper bounds on the price of anarchy for games defined on general network topologies; along the way we obtain an upper bound of 2 for games defined on star networks. We next show that our bounds are tight even in the case of tree networks of maximum degree 3, leading to non-constant price of anarchy for such topologies. In contrast, for network topologies of maximum degree 2 the quality of the solutions obtained by selfish wavelength assignment is much more satisfactory: we prove that the price of anarchy is bounded by 4 for a large class of practically interesting games defined on ring networks.

***Index Terms*—selfish wavelength assignment, non-cooperative games, bottleneck games, price of anarchy, multifiber optical networks, path multicoloring.**

I. INTRODUCTION

CENTRALIZED decision making in contemporary large-scale computer networks is often impractical or infeasible. Indeed, the relevant resource allocation problems usually turn out to be computationally intractable, thus forcing network operators to content themselves with suboptimal solutions that can be produced at a more realistic computational cost. In this light, a trend that constantly gains ground is to study the effect of reducing or even completely abandoning centralized resource allocation to network users. Such decentralized systems have been proposed and studied in the past in the context of routing. However, in the last decade there has been renewed interest in applying game-theoretic techniques to a variety of problems in networking. For numerous recent and earlier results on the subject the reader is referred to [1]–[3] and to references therein.

In this work, we draw motivation from large-scale optical networks which are currently deployed to sustain the bulk of network traffic generated by the ever-increasing and diversifying user demands. Unfortunately, the majority of resource allocation problems in optical networks are computationally intractable, even in restricted network topologies. We adopt an increasingly popular approach to modeling the lack of centralized control in optical networks, viewing the users as active, selfish, non-cooperating entities who compete with one another over network resources ([4]–[7]). We study the corresponding strategic game in a game-theoretic framework pioneered by Koutsoupias and Papadimitriou ([8], [9]) and we provide, among other results, a complete analysis of the deterioration of the quality of solutions caused by the lack of user coordination in a natural model for non-cooperative wavelength assignment in multifiber optical networks.

The tremendous bandwidth available in optical fibers is most efficiently exploited through the use of *Wavelength Division Multiplexing (WDM)*. WDM allows for splitting the fiber bandwidth into multiple independent channels (wavelengths), each one operating at a different light frequency. It is highly desirable that all communication in an optical network should be carried out *transparently*, i.e., each signal should use the same wavelength on all the fibers that it traverses. Transparency allows the use of all-optical switching components, and thus obviates the need for opto-electronic conversion that increases the cost of terminal equipment and slows down the network.

In a WDM network, communication requests that are routed on the same fiber are not allowed to use the same wavelength. The bandwidth supported by currently deployed implementations of WDM amounts to at most two hundred wavelengths per fiber, and this number is not expected to change drastically in the near future. An immediate remedy to this situation is to implement physical links using multiple parallel fibers. Naturally, this boosts the available bandwidth and also allows for several requests that use the same wavelength to be routed on the same physical link, provided that each one uses a different fiber. Several optimization problems have been defined and studied in the multifiber setting (see the subsection on related work below). Our focus here is on pre-routed requests. Pre-routed requests arise in settings where routing is unique (acyclic network topologies), or when the path on which a particular request will be routed is decided

This work has been funded by the project "Thales," cofinanced by the European Social Fund (80%) and national resources (20%) in the framework of the Operational Programme "Education and life long learning." A preliminary version of this work appears in the Proceedings of the 19th International Symposium on Algorithms and Computation (ISAAC), LNCS vol. 5369, pp. 159–170, Springer, 2008.

E. Bampas, A. Pagourtzis, and K. Potika are with the School of Electrical & Computer Engineering, National Technical University of Athens, Greece. e-mails: {ebamp, pagour, epotik}@cs.ntua.gr

G. Pierrakos is with the department of Electrical Engineering & Computer Sciences, UC Berkeley, CA, USA. e-mail: georgios@cs.berkeley.edu

independently of the wavelength assignment procedure. This is the case when there are specific routing requirements, such as shortest-path routing, or when the lightpaths are set up in an earlier stage of the virtual topology design process.

In this work, we introduce and study a model where there is no central authority assigning wavelengths to communication requests. Instead, each request is considered as a player in a strategic game, in which every possible choice of frequency incurs some cost on the request, depending on the choices of other players as well. Specifically, the cost of a player that has chosen a particular frequency f is equal to the maximum number of requests using the same frequency f encountered on some edge along its path. Each player/request behaves selfishly, trying to choose a wavelength that will minimize her cost. We call these SELFISH PATH MULTICOLORING (S-PMC) games[1]. In game theory, numerous ways have been proposed to model the outcome of strategic games, by far the most popular approach being to assume that the game settles in some *Nash equilibrium*, a stable state of the game in which no player can reduce her cost by changing strategy unilaterally. The famous result of Nash [11] guarantees the existence of *mixed* equilibria, in which each player chooses each of her available strategies with some probability. Here, we will focus exclusively on *pure* equilibria, where each player plays exactly one of her strategies with probability 1. Early results on the properties of (pure) Nash equilibria of non-cooperative network routing games, as well as further motivation for the game-theoretic viewpoint, can be found in the work of Orda et al. [12]. From the point of view of the network operator, the minimization of the cost of optical fibers that are required to accommodate all communication requests is a compelling desideratum. Thus, it is in the operator's best interest if the players reach a "good" Nash equilibrium, in the sense that the cost of optical fibers required to accommodate the wavelength choices of the players is small. In S-PMC games, we measure the quality of a game state by using a *social cost* function that is equal to the maximum number of requests that use the same wavelength on the same edge.

The above game-theoretic model admits at least two different interpretations of the player cost and social cost functions. In the first scenario, the player cost function is primarily regarded as a simple *charging mechanism* employed by the network operator in order to steer users towards Nash equilibria that are beneficial for the system. Accordingly, the social cost function represents the maximum number of parallel optical fibers that will be required on some physical link of the network in order to accommodate all requests. This is a justified optimization target from the network operator's perspective. Indeed, in certain cases the specific fiber cost per link may be unknown or vary with time, and therefore a reasonable objective is to minimize the maximum fiber usage over all links of the network [13]. Furthermore, scenarios where there is a fixed number of fibers per link arise naturally in practice and have been considered in several studies ([14]–[17]); in such cases, the minimization of maximum fiber multiplicity reveals which traffic patterns can be served by an existing infrastructure.

[1] In the context of multifiber optical networks, the term *multicoloring* refers to the fact that we allow multiple requests that share an edge to receive the same color [10].

In the second scenario, the player cost is viewed rather as a disutility naturally suffered by the request due to the conditions currently present in the network. For example, users may prefer wavelengths that are not used in too many other fibers in order to leave the highest possible margin for transferring to another fiber in case of fiber failure or other "kick-out" events (e.g. due to the arrival of higher priority requests). The social cost is now interpreted as the maximum disutility suffered by any player in the network, a measure known as *egalitarian* social cost which has been frequently used for evaluating the quality of strategy profiles in the selfish resource allocation literature ([18]–[21], see also [3, chapter 17]).

Note, additionally, that S-PMC games can model selfish wavelength assignment in single-fiber optical networks. In this setting, a color multiplicity on an edge represents the congestion encountered by requests using the corresponding wavelength on that edge. A player may naturally seek to minimize the maximum wavelength congestion along her path, in order to maximize her throughput. Similar to the second scenario above, the social cost represents the maximum disutility of any player.

In the Koutsoupias-Papadimitriou framework, the loss incurred by the lack of centralized control is measured by the *price of anarchy*: the ratio of the social cost of the worst-case Nash equilibrium to the social cost of an optimal, centrally computed, strategy profile. A small price of anarchy implies that one can let the players play the game selfishly and converge to some Nash equilibrium, the cost of which will not be far from the optimal. A second measure of interest, introduced by Anshelevich et al. [22], is the *price of stability*: the ratio of the social cost of a best-case Nash equilibrium to the cost of an optimal solution. A small price of stability implies that there exists a strategy profile that may be *imposed* on the players, from which no player will have incentive to deviate, and the cost of which will not be far from the optimal.

A. Related Work

S-PMC games naturally correspond to the following optimization problem: given a multifiber network, a set of pre-routed requests (fixed routing), and a number of available wavelengths, find a wavelength assignment to minimize the maximum fiber requirement on any network link (the number of fibers required on a link is equal to the maximum wavelength multiplicity on that link). This problem has been studied in [23] under the name MIN-MAXFIBER-FIXEDROUTE, where they present a randomized algorithm for general graphs which achieves a logarithmic approximation ratio; a very similar optimization problem was studied in [24], where logarithmic approximation hardness was shown which applies also to MIN-MAXFIBER-FIXEDROUTE as observed in [23]. As regards specific topologies, the algorithms proposed in [10] for a related optimization problem directly give exact solutions for chain networks and 2-approximate solutions for ring and star networks for MIN-MAXFIBER-FIXEDROUTE. Similar bounds hold also for the variant of the problem where requests are not pre-routed (flexible routing).

Another optimization problem that has been considered in multifiber networks is the minimization of the total fiber requirement. Exact or constant-ratio approximation algorithms for this problem were given for chains, rings (with both fixed and flexible routing), and stars in [10], for trees in [25], and for spiders in [26]. For general graphs, logarithmic approximation algorithms and hardness results were given in [23] and [27], assuming fixed as well as flexible routing. Yet another variant seeks to minimize the number of wavelengths used, given a fixed number of fibers on each link. This variant was studied in [14], [15], [25], and [28], where constant ratio approximation algorithms were presented for various simple topologies.

Selfish path multicoloring games as such have not been considered before in the literature. Even in the single fiber setting, selfish path coloring has only recently been studied in [4]–[7]. Bilò and Moscardelli [4] present payment strategies that induce games possessing pure Nash equilibria. Later, Bilò et al. [5] extend these games by introducing information levels to the local knowledge of the players for computing their payments and give bounds for the price of anarchy in chains, rings and trees. In the work of Georgakopoulos et al. [6], additionally to providing results for the existence and computation of pure Nash equilibria of selfish routing and path coloring games, they also consider the complexity of recognizing and computing better Nash equilibria for such games under various payment functions. Milis et al. [7] present upper and, for the first time, lower bounds on the price of anarchy of selfish routing and path coloring games under cost functions that charge each player taking into account only her own strategy choice. Among other results, they show a constant price of anarchy for selfish path coloring games in rings.

Selfish path multicoloring games are closely related to bottleneck games ([20], [21]), a variation of congestion games where a player's cost is determined by her *maximum* latency instead of the usual cost which is the *sum* of her latencies.

In [20], Busch and Magdon-Ismail 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 pure Nash equilibrium (hence the price of stability is 1) and that the price of anarchy of the game is determined by topological properties of the network. In particular, they show that the price of anarchy 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 in the network.

A further generalization is the model of Banner and Orda [21], 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 price of anarchy for these games is unbounded.

Both models are more general than the model considered in this article, since a selfish path multicoloring game can be seen as a traffic routing game in a multigraph, where each edge is replaced with w parallel edges, one for each available color. Each player's strategy set then consists of w edge-disjoint source-destination paths, corresponding to the w available colors in the selfish path multicoloring model. However, our model fits better into the framework of optical networks for which we provide smaller upper bounds on the price of anarchy compared to the ones obtained in [20] and [21], as well as a better convergence rate to Nash equilibria. Moreover, our lower bounds naturally hold in the more general models of [20] and [21], but are not directly comparable to the lower bounds presented there.

The type of bottleneck games discussed above are usually called *network* bottleneck games to stress the fact that the set of resources available to the players are edges of a graph; another type are *general* bottleneck games, where there is no such combinatorial structure underlying the set of resources; moreover a bottleneck game is said to be *asymmetric* when players are allowed to have different strategy sets (in our setting this translates to different source and destination nodes). In light of the observation in the previous paragraph, our model can be seen as a special case of (asymmetric) network bottleneck games, with linear delay functions on the edges. In [29] the authors show that the problem of computing (any) α-approximate Nash equilibrium for general asymmetric bottleneck games in which the delay functions on the resources have bounded jumps is a PLS-complete problem (i.e., the natural, exponential in the worst case, sequence of best-response dynamics is probably the best we can hope for). Unfortunately, their result does not immediately imply anything about the complexity of computing a Nash equilibrium in our model, since their model is more general both in the strategy space (general vs network game) and in the type of delay functions allowed (bounded-jump vs linear). On the other hand, for the related setting of congestion games, it is shown in [30] that computing a Nash equilibrium is PLS-complete even in (asymmetric) network congestion games with linear delay functions on the edges.

In [29] the authors also provide an algorithm for computing a *strong* equilibrium (a generalization of Nash equilibrium) in matroid bottleneck games. In Section IV we study selfish path multicoloring games on a specific network topology we call "rooted-tree" topology and we provide an algorithm for computing a Nash equilibrium. It is easy to see that this is a special case of matroid bottleneck games and therefore the algorithm of [29] works in our setting as well, however our algorithm is faster and has the property that it always computes a Nash equilibrium of optimal social cost.

B. Contributions

In this paper we propose SELFISH PATH MULTICOLORING games as a model for studying the behavior of multifiber optical networks in which the wavelength assignment process is carried out by the users in a selfish manner. We present a thorough analysis of the price of anarchy of these games, as well as results on the existence and computation of Nash equilibria and on the rate of convergence to Nash equilibria.

We first observe that all SELFISH PATH MULTICOLORING games converge to some Nash equilibrium in a finite number of steps. We show a lower rate of convergence than the one known for the more general games of [20] and [21]. We also observe that there always exist Nash equilibria of optimal social cost: the price of stability is 1. The problem of finding such equilibria is, in general, computationally intractable, since the corresponding optimization problem is NP-hard. In this work, we are able to pinpoint a subclass of games defined on trees ("rooted-tree" topology), for which we provide a polynomial-time algorithm that computes optimal Nash equilibria. We then prove that for all games defined on stars, we can use an approximation algorithm for the corresponding optimization problem in order to compute approximate equilibria.

Our main contribution is the analysis of the price of anarchy of several classes of SELFISH PATH MULTICOLORING games. We first provide upper bounds for games defined on general graphs; in particular, we prove that on any graph, the price of anarchy is bounded by the length of the shortest maximum-cost request in any Nash equilibrium. This bound immediately yields a constant price of anarchy, namely 2, for the special case of star topologies. We next show that our bounds are tight even in the case of tree networks of maximum degree 3, yielding non-constant price of anarchy for such (and more general) topologies. In contrast, for networks of maximum degree 2, that is rings and chains, the situation is much more gratifying: we show that the price of anarchy is bounded by $\max\{4, \frac{w^2}{L}\}$, where L is the maximum load of the network and w is the number of available wavelengths per fiber. This means that for a large class of practically interesting games, the price of anarchy is at most 4.

II. PRELIMINARIES

Given an undirected graph $G=(V,E)$, a nonempty set $P=\{p_1,\dots,p_N\}$ of simple paths defined on G, and a nonempty 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. $L=\max_{e\in E} L(e)$. We will occasionally view a path $p\in P$ as a set of edges, therefore the notation $e\in p$ will mean that path p uses edge e and $|p|$ will stand for the length of path p.

A *coloring*, i.e. an assignment of colors to paths, will be denoted by a vector $\mathbf{c}=(c_1,\dots,c_N)$ in W^N: each coordinate c_i denotes the color assigned to path p_i. With respect to a coloring $\mathbf{c}$, we will make use of the following notation:

Definition 1 (Notation)**.**
1) $P^{(\mathbf{c})}(e,\alpha)$ will denote the set of paths that use edge e and are colored with color α.
2) $\mu^{(\mathbf{c})}(e,\alpha)$ will denote the *multiplicity* of color α on edge e:
$$\mu^{(\mathbf{c})}(e,\alpha)=\left|P^{(\mathbf{c})}(e,\alpha)\right| \ .$$
3) $\mu_e^{(\mathbf{c})}$ will denote the maximum multiplicity of any color on edge e:
$$\mu_e^{(\mathbf{c})}=\max_{\alpha\in W}\mu^{(\mathbf{c})}(e,\alpha) \ .$$
4) $\mu_{\max}^{(\mathbf{c})}$ will denote the maximum multiplicity of any color over all edges:
$$\mu_{\max}^{(\mathbf{c})}=\max_{e\in E}\mu_e^{(\mathbf{c})} \ .$$
5) $\mu^{(\mathbf{c})}(p,\alpha)$ will denote the maximum multiplicity of color α over the edges of path p:
$$\mu^{(\mathbf{c})}(p,\alpha)=\max_{e\in p}\mu^{(\mathbf{c})}(e,\alpha) \ .$$

Whenever there is no ambiguity regarding the coloring we are referring to, we will not make explicit the dependence on $\mathbf{c}$.

Note that $\mu_{\max}^{(\mathbf{c})}$, as defined above, represents the cost of a coloring in the corresponding optimization problem. The minimum cost over all possible colorings $\mathbf{c}$ will be denoted by μ_{OPT}, i.e.:
$$\mu_{\mathrm{OPT}}=\min_{\mathbf{c}}\mu_{\max}^{(\mathbf{c})} \ ,$$
where $\mathbf{c}$ ranges over all possible colorings. We observe immediately that in any coloring, at least one color will appear with multiplicity at least $\lceil\frac{L}{w}\rceil$ on the maximum-load edge. Therefore, we obtain the following lower bound on μ_{OPT}:

Fact 1. *No coloring can achieve a cost that is smaller than* $\lceil\frac{L}{w}\rceil$. *Thus,*
$$\mu_{\mathrm{OPT}}\geq\left\lceil\frac{L}{w}\right\rceil \ .$$

A. Game-Theoretic Model

We now define formally our game-theoretic model for non-cooperative (or selfish) wavelength assignment in multifiber optical networks.

Definition 2 (SELFISH PATH MULTICOLORING games)**.** Let G be an undirected graph, $P=\{p_1,\dots,p_N\}$ be a set of simple paths defined on G, and $W=\{\alpha_1,\dots,\alpha_w\}$ be a set of available colors. The SELFISH PATH MULTICOLORING game $\langle G,P,W\rangle$ is defined as follows:

- *Players*: there is one player for each path in P. For simplicity, we will identify each player i with the corresponding path p_i.
- *Strategies*: All players share the common set of available strategies W. The choice of strategy (color) of player i is denoted by $c_i\in W$. A *strategy profile* for the game is a coloring $\mathbf{c}=(c_1,\dots,c_N)$ that corresponds to the color choices made by the players.
- *Disutility*: The *disutility* of player i in the strategy profile $\mathbf{c}$ is given by the disutility function $f_i: W^N\to\mathbb{N}$ as follows:
$$f_i(\mathbf{c})=\mu^{(\mathbf{c})}(p_i,c_i) \ .$$

Definition 3. S-PMC will denote the class of all SELFISH PATH MULTICOLORING games.

We will use the notation S-PMC(X) to denote a subclass of S-PMC that contains only games satisfying a property X. For example, the restriction of S-PMC to games defined on trees (resp. stars) is denoted by S-PMC(TREE) (resp. S-PMC(STAR)), where we use TREE for the property "G is a tree".

It is clear from Definition 2 that we concentrate on pure strategies and do not consider the case where players might pick each color with some probability. Following the standard definition, we say that a strategy profile $\mathbf{c}=(c_1,\dots,c_N)$ is a *pure Nash equilibrium* (PNE), or simply for our purposes *Nash equilibrium* (NE), if for each player i it holds that:

$$f_i(c_1, \ldots, c_i', \ldots, c_N) \geq f_i(c_1, \ldots, c_i, \ldots, c_N) \ ,$$

for any strategy $c_i' \in W$. Moreover, following the definition from Chien and Sinclair [31], we say that a strategy profile $\mathbf{c} = (c_1, \ldots, c_N)$ is an *ε-approximate Nash equilibrium* if for each player i it holds that:

$$f_i(c_1, \ldots, c_i', \ldots, c_N) \geq (1-\varepsilon) \cdot f_i(c_1, \ldots, c_i, \ldots, c_N) \ ,$$

for any strategy $c_i' \in W$. In a Nash equilibrium, no player will improve her disutility by changing strategy unilaterally, while in an ε-approximate Nash equilibrium a unilateral change of strategy may result in reducing the deviating player's cost by no more than a factor of $1-\varepsilon$.

Definition 4 (Blocking edges)**.** Let $\mathbf{c}$ be a strategy profile for some game $\langle G, P, W \rangle$. We say that edge e is an *α-blocking edge for $p_i \in P$*, or that it *blocks α for p_i*, if $e \in p_i$ and $\mu^{(\mathbf{c})}(e, \alpha) \geq f_i(\mathbf{c}) - 1$. Furthermore, in that case, the paths in $P^{(\mathbf{c})}(e, \alpha)$ are called *α-blocking paths for p_i*.

Intuitively, an α-blocking edge for p_i "blocks" this player from switching to color α because if she did, her new disutility would be at least $\mu^{(\mathbf{c})}(e, \alpha) + 1 \geq f_i(\mathbf{c})$, no better than her disutility in the current coloring. The following characterization of the Nash equilibria of S-PMC games is immediate from the definitions:

Property 2 (Structural characterization of S-PMC Nash equilibria)**.** *A strategy profile for an* S-PMC *game $\langle G, P, W \rangle$ is a Nash equilibrium if and only if every path $p \in P$ contains at least one α-blocking edge for p, for every color α.*

Definition 5 (Social cost)**.** We define the *social cost* of a strategy profile $\mathbf{c}$ for an S-PMC game to be the cost of the corresponding coloring:

$$\mathrm{sc}(\mathbf{c}) = \mu_{\max}^{(\mathbf{c})} \ .$$

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

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

We define $\hat{\mu}$ to be the maximum social cost over all strategy profiles that are Nash equilibria:

$$\hat{\mu} = \max_{\mathbf{c} \text{ is NE}} \mathrm{sc}(\mathbf{c}) \ .$$

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.:

$$\mathrm{PoA}(\langle G, P, W \rangle) = \frac{\max_{\mathbf{c} \text{ is NE}} \mathrm{sc}(\mathbf{c})}{\mu_{\mathrm{OPT}}} = \frac{\hat{\mu}}{\mu_{\mathrm{OPT}}} \ .$$

The *price of stability* (PoS) of a game is the best-case social cost in a Nash equilibrium divided by μ_{OPT}:

$$\mathrm{PoS}(\langle G, P, w \rangle) = \frac{\min_{\mathbf{c} \text{ is NE}} \mathrm{sc}(\mathbf{c})}{\mu_{\mathrm{OPT}}} \ .$$

The price of anarchy (resp. stability) of a class of games S-PMC(X) is the least upper bound on the price of anarchy (resp. stability) of S-PMC games that satisfy property X.

III. Price of Stability, Existence, and Convergence to Equilibria

We first prove that any S-PMC game $\langle G, P, W \rangle$ has at least one Nash equilibrium of optimal social cost. Moreover, we show that starting from an arbitrary strategy profile, any *Nash dynamics* converges to a Nash equilibrium in at most 4^N steps. For our purposes, the Nash dynamics is a sequence $\mathbf{c}_0, \mathbf{c}_1, \ldots$ of strategy profiles where in each profile $\mathbf{c}_{i+1}$ exactly one player has a different strategy compared to $\mathbf{c}_i$; moreover, that player has strictly decreased her disutility compared to her disutility in $\mathbf{c}_i$. In other words, the Nash dynamics is a sequence of cost-improving moves of the players in which no particular order of play or fairness criteria is assumed *a priori.*

For any strategy profile $\mathbf{c}$, we consider a disutility vector $\boldsymbol{D}(\mathbf{c})$ defined as follows:

$$\boldsymbol{D}(\mathbf{c}) = (d_L(\mathbf{c}), \ldots, d_1(\mathbf{c})) \ ,$$

where $d_i(\mathbf{c})$ stands for the number of players whose disutility is exactly i (note that the disutility of any player cannot be 0 and cannot be greater than L). We use lexicographic-order arguments similar to those in [20] and [21] to show that starting from an arbitrary strategy profile any Nash dynamics converges to a Nash equilibrium of smaller or equal social cost.

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

1) *the price of stability is 1, and*
2) *any Nash dynamics converges to a Nash equilibrium in at most 4^N steps.*

Proof: Let $\prec$ denote the standard lexicographic ordering between vectors of equal size. If $\mathbf{c}$ is a strategy profile for $\langle G, P, W \rangle$ that is not a Nash equilibrium and $\mathbf{c}'$ is the strategy profile resulting from a profitable deviation of some player p_i, we show that $\boldsymbol{D}(\mathbf{c}') \prec \boldsymbol{D}(\mathbf{c})$ and hence $\mathrm{sc}(\mathbf{c}') \leq \mathrm{sc}(\mathbf{c})$. This implies that any Nash dynamics starting from a minimum-cost strategy profile converges to a Nash equilibrium of the same social cost, thus the price of stability is 1.

To prove the claim, we show that all players whose disutility changes in $\mathbf{c}'$ have a new disutility strictly smaller than $f_i(\mathbf{c})$. This guarantees that the new disutility vector is lexicographically smaller than the previous one. Clearly, this holds for player p_i herself.

Some of the players that overlap with p_i and are colored with c_i may also have their disutilities reduced by exactly 1. The original disutility of any such player p_j must be $f_j(\mathbf{c}) \leq f_i(\mathbf{c})$, therefore $f_j(\mathbf{c}') \leq f_i(\mathbf{c}) - 1$. On the other hand, the deviation of player i may result in an increase by exactly 1 of the disutility of some players who overlap with p_i and are colored with c_i'. For any player p_k whose disutility is increased, it holds $f_k(\mathbf{c}) \leq f_i(\mathbf{c}) - 2$, otherwise p_i would be blocked from switching to p_k's color. Therefore, $f_k(\mathbf{c}') \leq f_i(\mathbf{c}) - 1$ and the claim is proved.

Regarding the rate of convergence, observe that for any strategy profile $\mathbf{c}$ the sum of the components of the corresponding disutility vector $\boldsymbol{D}(\mathbf{c})$ is:

$$\sum_{i=1}^{L} d_i(\mathbf{c}) = N \ ,$$

independent of $\mathbf{c}$. Therefore, the number of distinct disutility vectors is at most equal to the number of distinct ways in which N indistinguishable balls can be thrown into L bins. This number is equal to

$$\binom{N+L-1}{L-1} \leq 2^{N+L-1} < 4^N \ ,$$

because $L \leq N$. The convergence of any Nash dynamics in at most this many steps follows, since any cost-improving move results in a new disutility vector which is lexicographically strictly smaller than the current disutility vector. ■

IV. Computing Optimal and Approximate Equilibria

In view of Theorem 1, computing a Nash equilibrium of minimum social cost is at least as hard as the corresponding optimization problem, in which one is given a graph G, a set of simple paths P defined on G, and the number of available colors w and is asked to color all paths in P so that the maximum fiber multiplicity $\mu_{\max}$ is minimized. Using a simple reduction from the single fiber path coloring problem (as was done in [10] for a similar path multicoloring problem) it can be shown that the problem is NP-hard in general graphs, in fact even in rings and stars (see also [13]). 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 S-PMC(Tree). Furthermore, we show that we can use a known algorithm for the Path MultiColoring problem in stars to compute $\frac{1}{2}$-approximate Nash equilibria for games in S-PMC(Star).

Definition 6. We define S-PMC(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 root node r such that each path in P lies entirely on some simple path from r to a leaf."

A similar class of instances has been defined and studied as an intersection model for "rooted directed edge path graphs" in [32].

We will say that a path in a tree rooted at node r *starts* from edge e, if e is the edge of the path that lies closest to node r. The following algorithm is a polynomial-time algorithm that computes optimal Nash equilibria for S-PMC(Rooted-Tree) games. It greedily colors paths in order of non-decreasing distance of their starting edge in such a way that the disutility of the path at the time of coloring is the lowest possible with respect to the current partial coloring:

Algorithm Rooted-Tree-NE

Input: an S-PMC(Rooted-Tree) game $\langle G, P, W\rangle$

Output: an optimal Nash equilibrium for $\langle G, P, W\rangle$

Find a root node r such that each path in P lies on some simple path from r to a leaf.
for all edges $e \in E$ in order of non-decreasing distance from r, breaking ties arbitrarily **do**
  for all paths p that start from edge e **do**
    Pick a color α such that $\mu(e, \alpha)$ is minimum in the current coloring, breaking ties arbitrarily.
    Assign color α to path p.
  end for
end for

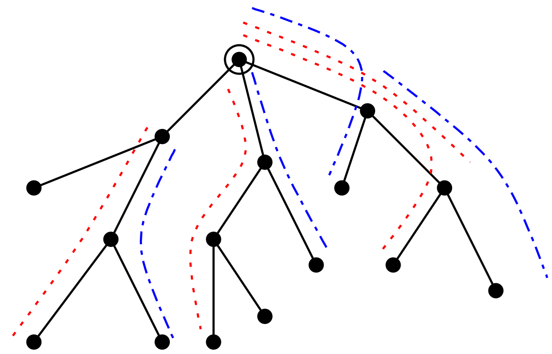

Fig. 1. A coloring obtained by Rooted-Tree-NE given an instance with $|W| = 2$. Different wavelengths are illustrated by different line style/color combinations. Solid black lines between nodes represent the edges of the underlying graph. The root node is marked with a small circle.

A coloring obtained by the algorithm is illustrated in Fig. 1. We proceed to prove its correctness.

Theorem 2. *Algorithm* Rooted-Tree-NE *computes a Nash equilibrium of optimal social cost* $\lceil \frac{L}{w} \rceil$ *for any* S-PMC(Rooted-Tree) *game.*

Proof: Let $e_1, \ldots, e_{|E|}$ be the order in which the algorithm considers the edges of G, let P_i be the subset of paths that are colored after the i-th iteration of the outer loop, and let $\mathbf{c}_i$ be the corresponding partial coloring. Note that P_1 contains exactly the paths that start from edge e_1 and, for all $i > 1$, $P_i \setminus P_{i-1}$ contains exactly the paths that start from edge e_i. Finally, note that $P_{|E|} = P$.

We first prove that the coloring returned by the algorithm is a Nash equilibrium. More specifically, we will show that, for all $i \geq 1$, the strategy profile $\mathbf{c}_i$ is a Nash equilibrium for the game $\langle G, P_i, W\rangle$. For $i = 1$, because the multiplicity of any color on edge e_1 is either $\left\lfloor \frac{L(e_1)}{w} \right\rfloor$ or $\left\lceil \frac{L(e_1)}{w} \right\rceil$ ($L(e_1)$ is the load on edge e_1 with respect to the complete path set P), no path in P_1 has incentive to change color and $\mathbf{c}_1$ is a Nash equilibrium for $\langle G, P_1, W\rangle$.

For the inductive step, assume that $\mathbf{c}_{i-1}$ is a Nash equilibrium for $\langle G, P_{i-1}, W\rangle$, for some $i > 1$. Let p be a path in P_i and let α be the color assigned to p in the profile $\mathbf{c}_i$. First, assume that $\mu^{(\mathbf{c}_i)}(p, \alpha) = \mu^{(\mathbf{c}_{i-1})}(p, \alpha)$, so that the disutility of p after the i-th iteration is exactly the same as it was after the $(i-1)$-st iteration. Moreover, since $P_i \supseteq P_{i-1}$, path p still contains after the i-th iteration at least the blocking edges that it contained after the $(i-1)$-st iteration. Therefore, path p has no incentive to change strategy in the profile $\mathbf{c}_i$.

Now assume that $\mu^{(\mathbf{c}_i)}(p, \alpha) > \mu^{(\mathbf{c}_{i-1})}(p, \alpha)$. Considering that only paths which start on edge e_i are assigned colors during the i-th iteration, this implies that p contains e_i and the maximum multiplicity of color α along path p in the profile $\mathbf{c}_i$

appears on edge e_i or on some edge e_j with $j > i$. However, any colored path that contains e_j also contains e_i, therefore:

$$\mu^{(\mathbf{c}_i)}(p,\alpha) = \mu^{(\mathbf{c}_i)}(e_i,\alpha) \ . \tag{1}$$

Let p' be the last path to be colored with color α during the i-th iteration, among the paths that start on edge e_i. At the moment p' was colored, color α must have been a minimum multiplicity color. Therefore, for any color $\alpha' \neq \alpha$:

$$\mu^{(\mathbf{c}_i)}(e_i,\alpha') \geq \mu^{(\mathbf{c}_i)}(e_i,\alpha) - 1 \ . \tag{2}$$

From (1) and (2), edge e_i is an α'-blocking edge for p, for any color $\alpha' \neq \alpha$, and thus p has no incentive to change color.

We have thus proved that the algorithm computes a Nash equilibrium. Note that, since $P_i \supseteq P_{i-1}$ for all i (assume that $P_0 = \emptyset$), $\mu_{\max}^{(\mathbf{c}_i)}$ is a non-decreasing function of i. Consider the last iteration j for which $\mu_{\max}^{(\mathbf{c}_j)} > \mu_{\max}^{(\mathbf{c}_{j-1})}$. Clearly, $\mu_{\max}^{(\mathbf{c}_j)}$ took its final value after the coloring of a certain path which starts from edge e_j with some color α. From (2) and the fact that α is a maximum-multiplicity color on e_j, it follows that color multiplicities on e_j are equal or differ at most by 1, hence

$$\mu_{\max}^{(\mathbf{c}_j)} = \mu_{e_j}^{(\mathbf{c}_j)} = \left\lceil \frac{L(e_j)}{w} \right\rceil \leq \left\lceil \frac{L}{w} \right\rceil \ .$$

Using Fact 1, it turns out that for the final coloring $\mathbf{c} = \mathbf{c}_{|E|}$, we have

$$\mu_{\max}^{(\mathbf{c})} = \mu_{\mathrm{OPT}} = \left\lceil \frac{L}{w} \right\rceil \ .$$

■

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

Proof: Let $\langle G,P,W\rangle$ be a game in S-PMC(STAR). We use the polynomial-time approximation algorithm presented by Nomikos et al. in [10] for the PATH MULTICOLORING problem in stars. This algorithm returns a coloring of the paths in P with the following property: for any edge e, the paths that use e can be partitioned into two sets of cardinality $L_1(e)$ and $L_2(e)$ respectively, such that for any color α

$$\left\lceil \frac{L_1(e)}{w} \right\rceil + \left\lceil \frac{L_2(e)}{w} \right\rceil - 2 \leq \mu(e,\alpha) \leq \left\lceil \frac{L_1(e)}{w} \right\rceil + \left\lceil \frac{L_2(e)}{w} \right\rceil \ .$$

Note that any player who changes color causes an increase by 1 of the multiplicity of the new color on the edges used by that player. Together with the above inequalities, this implies that if $\mathbf{c}$ is the strategy profile returned by the algorithm, then any player i who deviates resulting in a new strategy profile $\mathbf{c}'$ may reduce her cost by at most 1. Therefore,

$$f_i(\mathbf{c}') \geq f_i(\mathbf{c}) - 1 = \left(1 - \frac{1}{f_i(\mathbf{c})}\right) \cdot f_i(\mathbf{c}) \ . \tag{3}$$

Now, players with $f_i(\mathbf{c}) = 1$ certainly have no incentive to deviate in the profile $\mathbf{c}$, therefore in the worst case we have $f_i(\mathbf{c}) = 2$ in (3). Hence $\mathbf{c}$ is an $\frac{1}{2}$-approximate Nash equilibrium. ■

V. The Price of Anarchy in General Graphs

Although we can efficiently compute optimal or close-to-optimal Nash equilibria for certain classes of games, we show in this section that the quality of obtained solutions may deteriorate if the players are allowed to decide their own strategy. In other words, there exist families of S-PMC games in which the social cost of the worst-case Nash equilibrium is arbitrarily high compared to the optimum. In such games, the player-charging mechanism of our model is inadequate to steer players towards Nash equilibria that are beneficial for the network.

More specifically, for any S-PMC game we exhibit two upper bounds on the price of anarchy. 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 then prove that these bounds are tight 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 trees with maximum degree equal to 3.

The two upper bounds are stated in Lemmas 3 and 4 below.

Lemma 3. *The price of anarchy of any* S-PMC *game $\langle G,P,W\rangle$ is at most w.*

Proof: The cost $\hat{\mu}$ of the worst-case Nash equilibrium of $\langle G,P,W\rangle$ cannot exceed L, so Fact 1 yields the following bound: $\mu_{\mathrm{OPT}} \geq \lceil \frac{L}{w} \rceil \geq \frac{L}{w} \geq \frac{\hat{\mu}}{w}$. Therefore, $\mathrm{PoA}(\langle G,P,W\rangle) = \frac{\hat{\mu}}{\mu_{\mathrm{OPT}}} \leq w$. ■

Lemma 4. *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 $f_i(\mathbf{c}) = \mathrm{sc}(\mathbf{c}) = \hat{\mu}$, the price of anarchy of game $\langle G,P,W\rangle$ is at most equal to the length of path p_i.*

Proof: Let z denote the length of path p_i. Since $\mathbf{c}$ is a Nash equilibrium, by Property 2, for each color $\alpha \in W$ path p_i must contain at least one α-blocking edge for p_i. Thus, there must be some edge $e \in p_i$ that blocks at least $\lceil \frac{w}{z} \rceil$ distinct colors for p_i. Since the disutility of path p_i is $f_i(\mathbf{c}) = \hat{\mu}$, the load of edge e is at least:

$$L(e) \geq 1 + \left\lceil \frac{w}{z} \right\rceil \cdot (\hat{\mu} - 1) \ .$$

By Fact 1 and the above inequality, we get that

$$\mu_{\mathrm{OPT}} \geq \left\lceil \frac{L}{w} \right\rceil \geq \left\lceil \frac{L(e)}{w} \right\rceil \geq \left\lceil \frac{1 + \lceil \frac{w}{z} \rceil \cdot (\hat{\mu} - 1)}{w} \right\rceil \ .$$

Therefore, the price of anarchy is bounded as follows:

$$\mathrm{PoA}(\langle G,P,W\rangle) = \frac{\hat{\mu}}{\mu_{\mathrm{OPT}}} \leq \frac{\hat{\mu}}{\left\lceil \frac{1 + \lceil \frac{w}{z} \rceil \cdot (\hat{\mu}-1)}{w} \right\rceil} \ . \tag{4}$$

Now, let $\hat{\mu} = \lambda z + \chi$, where λ and χ are integers satisfying

$\lambda \geq 0$ and $0 \leq \chi \leq z-1$. We can rewrite (4) as follows:

$$\begin{aligned}\text{PoA}(\langle G,P,W\rangle) &\leq \frac{\lambda z+\chi}{\left\lceil \left\lceil \frac{w}{z}\right\rceil \cdot \frac{z}{w}\cdot\lambda + \frac{\left\lceil \frac{w}{z}\right\rceil\cdot(\chi-1)+1}{w}\right\rceil} \\ &\leq \frac{\lambda z+\chi}{\left\lceil \lambda + \frac{\left\lceil \frac{w}{z}\right\rceil\cdot(\chi-1)+1}{w}\right\rceil}\ , \end{aligned} \quad (5)$$

because $\left\lceil \frac{w}{z}\right\rceil \geq \frac{w}{z}$. Now if $\chi = 0$, (5) gives:

$$\text{PoA}(\langle G,P,W\rangle) \leq \frac{\lambda z}{\left\lceil \lambda - \frac{\left\lceil \frac{w}{z}\right\rceil -1}{w}\right\rceil} \leq \frac{\lambda z}{\left\lceil \lambda - \frac{w-1}{w}\right\rceil}\ ,$$

because $\left\lceil \frac{w}{z}\right\rceil \leq w$. By the last inequality we get $\text{PoA}(\langle G,P,W\rangle) \leq \frac{\lambda z}{\lambda} = z$.

On the other hand, if $1 \leq \chi \leq z-1$, (5) gives:

$$\text{PoA}(\langle G,P,W\rangle) \leq \frac{\lambda z+z-1}{\left\lceil \lambda + \frac{1}{w}\right\rceil} = \frac{\lambda z+z-1}{\lambda+1} < z\ .$$

■

As an immediate corollary of Lemma 4, we derive the following upper bound on the price of anarchy:

Corollary 5. *The price of anarchy of any* S-PMC *game* $\langle G,P,W\rangle$ *is bounded as follows:*

$$\text{PoA}(\langle G,P,W\rangle) \leq \min_{\mathbf{c}\text{ is NE}:\text{sc}(\mathbf{c})=\hat{\mu}}\ \min_{i:f_i(\mathbf{c})=\hat{\mu}} |p_i|\ .$$

In fact, the upper bounds stated in Lemma 3 and Corollary 5 are tight for rooted-tree games and asymptotically tight for rooted-tree games of maximum degree 3. In Lemmas 6 and 7 below, we describe the construction of families of games which exhibit a price of anarchy that matches these upper bounds.

Lemma 6. *The upper bounds of Lemma 3 and Corollary 5 are tight 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 Fig. 2.

For any $z \geq 1$ and $\lambda \geq 1$, let $\mathcal{A}_z(\lambda)$ be the following S-PMC game with z available colors: there are λ paths of color α_1 and length z which branch out into λ branches, one on each branch. Let us call these the *primary* paths for $\mathcal{A}_z(\lambda)$. 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 $\mathcal{A}_z(\lambda-1)$ game. They become primary paths for this copy of $\mathcal{A}_z(\lambda-1)$. The base case of this recursive construction is $\mathcal{A}_z(0)$, which is a degenerate game with no paths, defined on a graph consisting of a single node. We have included the explicit construction for $z=\lambda=3$ in Fig. 3.

Claim. For any $z \geq 1$, the construction $\mathcal{A}_z(z)$ is an S-PMC(ROOTED-TREE) game in Nash equilibrium, in which all of the following are equal to z: the number of available colors w, the maximum load L, the maximum color multiplicity $\mu_{\max}$, and all path lengths.

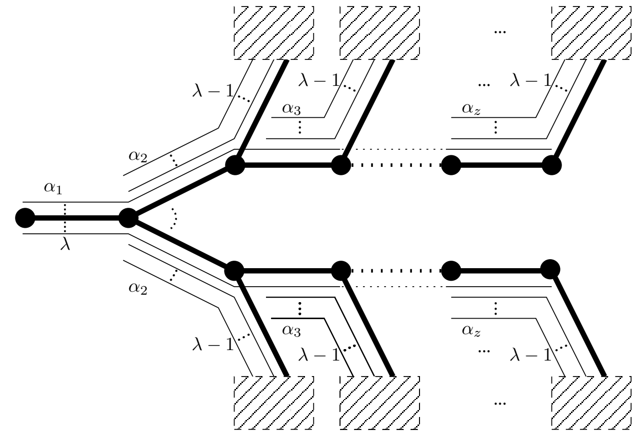

Fig. 2. The construction $\mathcal{A}_z(\lambda)$ for the proof of Lemma 6. The thick lines represent the edges of the underlying graph, and the thin lines represent the paths defined on the graph. The color and multiplicity of each group of paths that use the same edge and have the same color are displayed next to that group. Each shaded box represents a recursive copy of $\mathcal{A}_z(\lambda-1)$.

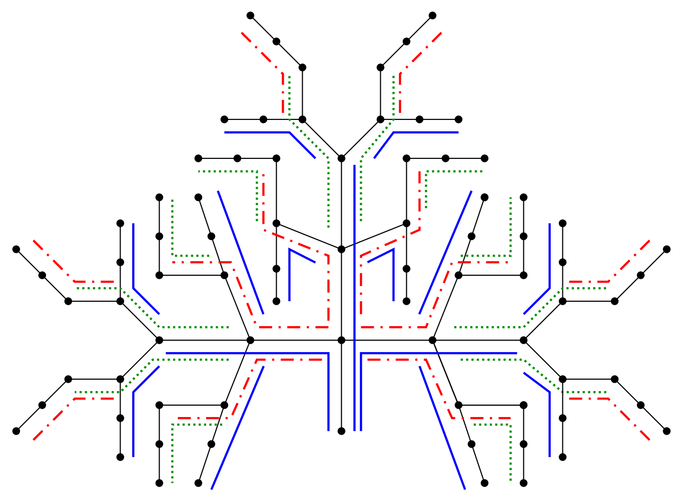

Fig. 3. The construction $\mathcal{A}_3(3)$, as described in the proof of Lemma 6. Different wavelengths are illustrated by different line style/color combinations. Solid black lines between nodes represent the edges of the underlying graph.

Proof of Claim: It is straightforward to verify that $\mathcal{A}_z(z) \in$ S-PMC(ROOTED-TREE); the root node is the root node u_0 of the first level of the recursive construction. The game is in Nash equilibrium by construction, since every path contains one blocking edge for every color other than its own. The number of available colors is equal to z by definition. The maximum multiplicity $\mu_{\max}=z$ appears on the edge incident to the root node of $\mathcal{A}_z(z)$. The maximum load $L=z$ appears on all the edges of the first level of the construction. Finally, all path lengths are equal to z by construction. The claim is proved.

By Theorem 2, the optimal strategy profile for $\mathcal{A}_z(z)$ has social cost $\mu_{\text{OPT}} = \left\lceil \frac{L}{w}\right\rceil = 1$. Therefore, the ratio $\frac{\mu_{\max}}{\mu_{\text{OPT}}}$ is equal to z for the Nash equilibrium we constructed, hence the price of anarchy is at least z. ■

Lemma 7. *The upper bounds of Lemma 3 and Corollary 5 are asymptotically tight for the class of* S-PMC(ROOTED-TREE: $\Delta=3$) *games.*

Proof: The construction presented in Lemma 6 can be modified so that the maximum degree of the resulting tree is 3, with only a logarithmic increase in the length of the paths.

For $z \geq \lambda$, we define the new construction $\mathcal{A}'_z(\lambda)$ of an S-PMC(ROOTED-TREE: $\Delta = 3$) game and Nash equilibrium thereof as follows: the underlying graph is the same as the underlying graph of $\mathcal{A}_z(\lambda)$, except that in any recursive copy of $\mathcal{A}'_z(\lambda-1)$, we interject between nodes u_1 and $u_{1,2}, \ldots, u_{\lambda,2}$ (cf. Fig. 2) a single edge out of node u_1 followed by a binary tree with height $\lceil \log z \rceil \geq \lceil \log \lambda \rceil$ and exactly λ leaves, which coincide with the nodes $u_{1,2}, \ldots, u_{\lambda,2}$. The extra edge interjected before the binary tree is necessary to ensure degree 3 in the later stages of the recursive construction; if we do not include this edge, but rather branch out into the first level of the binary tree, then node $u_{1,2}$ for example will have degree four. The set of primary paths of each recursive copy of $\mathcal{A}'_z(\lambda - 1)$ also remains the same, except that for all i: $1 \leq i \leq \lambda$, the primary path that used edge $(u_1, u_{i,2})$ in $\mathcal{A}_z(\lambda)$ is now stretched to use the edges connecting u_1 to $u_{i,2}$ through the newly interjected binary tree. Therefore, all paths now have the same length $\ell = z + \lceil \log z \rceil + 1$. The number of available colors and the coloring of the paths is the same as in the original construction.

This process results in an underlying graph which is a tree of maximum degree 3. The strategy profile is also a Nash equilibrium for the new game, since there is no change in the overlaps between paths. The construction $\mathcal{A}'_z(z)$ is an S-PMC(ROOTED-TREE: $\Delta = 3$) game in Nash equilibrium, with the same properties as the construction in Lemma 6 except that the length of all paths is exactly $\ell = z + \lceil \log z \rceil + 1 = \Theta(z)$. It turns out, then, that $\mathrm{PoA} \geq z = \ell - (\lceil \log z \rceil + 1) = \ell - o(z) = \ell - o(\ell)$. ∎

We summarize the results of Lemmas 3, 4, 6, and 7 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*

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

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

On a final note, we prove a constant, tight bound of 2 on the price of anarchy of S-PMC games defined on stars.

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

Proof: Lemma 4 implies an upper bound of 2 on the price of anarchy, since the length of any simple path defined on a star cannot be greater than 2.

For the lower bound, we can easily modify the construction that appears in the proof of Lemma 6 to yield a family of S-PMC(STAR) games with price of anarchy 2. More specifically, observe that any game $\mathcal{A}_2(\lambda)$ contains only players (paths) of length 2. Such a game can be embedded in a star with exactly the same number of edges as follows: fix an isomorphism φ between the edges of the tree and the edges of the star, and for every player $p = \{e, e'\}$ defined on the tree, define a player $\tilde{p} = \{\varphi(e), \varphi(e')\}$ with the same color on the star. It is clear that the paths we just defined on the star overlap with each other in exactly the same way as the original paths overlapped on the tree. Therefore, the game on the star is in Nash equilibrium with $\mu_{\max} = \lambda$, whereas the optimal solution has maximum color multiplicity $\mu_{\mathrm{OPT}} = \lceil \frac{L}{w} \rceil = \lceil \frac{\lambda}{2} \rceil$, where we used the fact that we have a special case of rooted trees and so $\mu_{\mathrm{OPT}} = \lceil \frac{L}{w} \rceil$ according to Theorem 2. ∎

VI. THE PRICE OF ANARCHY IN GRAPHS OF MAXIMUM DEGREE 2

In the previous section we gave two generic upper bounds on the price of anarchy of S-PMC games on general graphs, and provided matching lower bounds for graphs of maximum degree at least 3. We now proceed to determine the price of anarchy of this model on graphs of maximum degree 2. This graph family contains the fundamental network topology of rings, thus it is interesting to study the price of anarchy of the classes S-PMC(RING) and S-PMC(CHAIN) from a theoretical as well as from a practical point of view.

We start by proving, in Lemma 8, a stronger necessary condition for Nash equilibria of S-PMC(RING) games, compared to the one we have already seen in Property 2 for Nash equilibria of arbitrary S-PMC games. Then, we employ this structural property in order to show, in Lemma 9, that any S-PMC(RING) game with $\hat{\mu} \geq w$ necessarily contains an edge with high load. This allows us to prove a constant upper bound on the price of anarchy for a broad class of S-PMC(RING) games with $L = \Omega(w^2)$. We will refer to these games as "heavily loaded." 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 the maximum load can be arbitrarily large depending on network traffic. Finally, for the sake of completeness, we show that the price of anarchy may become unbounded if the network designer opts to provide the users with ample wavelengths, i.e. when $L = o(w^2)$. Instances with $L = o(w^2)$ will be called "lightly loaded."

We introduce the following notation: For an S-PMC(RING) game, $[e, e']$ denotes the arc of the ring that contains all edges between e and e' in the clockwise direction, including e and e'. Additionally, we say that $[e, e']$ is *contained* in a path p if every edge in $[e, e']$ belongs to p.

Lemma 8 (Structural property of S-PMC(RING) Nash equilibria)**.** *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 *there is an arc* $[e^-, e^+]$ *(that contains* e*) with the following properties:*

1) *for every edge* e' *of the arc* $[e^-, e^+]$ *it holds that*

$$\mu(e', \alpha) \geq |P(e', \alpha) \cap P(e, \alpha)| \geq \left\lceil \frac{\mu(e, \alpha)}{2} \right\rceil \ ,$$

and

2) *for every color* α'*, there is an edge* e' *of the arc* $[e^-, e^+]$ *such that*

$$\mu(e', \alpha') \geq \mu(e, \alpha) - 1 \ .$$

Proof: The strategy profile $\mathbf{c}$ is a Nash equilibrium, therefore by Property 2 every path must contain at least one α'-blocking edge, for every color α'. The paths in $P(e,\alpha)$ have cost at least $\mu(e,\alpha)$, therefore for every color α' there must exist at least one edge e_b on the ring with $\mu(e_b,\alpha') \geq \mu(e,\alpha)-1$.

Now, fix a color α' and consider the edge e_1 (resp. e_2) that lies closest to e in the clockwise (resp. counter-clockwise) direction and for which $\mu(e_1,\alpha') \geq \mu(e,\alpha)-1$ (resp. $\mu(e_2,\alpha') \geq \mu(e,\alpha)-1$). It may well be the case that e_1 is identical to e_2. We now observe that either $[e,e_1]$ is contained in at least half of the paths in $P(e,\alpha)$ or $[e_2,e]$ is contained in at least half of the paths in $P(e,\alpha)$, otherwise there would exist at least one path in $P(e,\alpha)$ which would not contain an α'-blocking edge. In the first case, we define $b(\alpha')=e_1$ and designate α' as a *positive* color, otherwise we define $b(\alpha')=e_2$ and designate α' as a *negative* color. Note that, by definition, a color cannot be both positive and negative at the same time.

We now define two sets B^+ and B^-: B^+ consists of the edges $b(\alpha')$ for all positive colors α', whereas B^- consists of the edges $b(\alpha')$ for all negative colors α'. Note that α is always a positive color and thus $b(\alpha)=e\in B^+$. Let e^+ be the last edge of B^+ encountered during a clockwise traversal of the ring starting at edge e, and let e^- be the first edge of B^- encountered during the same traversal (let $e^-=e$ if B^- is empty).

We prove that arc $[e^-,e^+]$ actually contains edge e. This is equivalent to stating that if B^- is not empty, then e^- is not contained in $[e,e^+]$. Indeed, if $e^-=b(\alpha')$ was contained in $[e,e^+]$, for some negative color α', then by the definition of edge e^+, arc $[e,e^-]$ would be contained in at least half of the paths in $P(e,\alpha)$ and therefore α' would be a positive and negative color at the same time, a contradiction.

Finally, we claim that the two properties in the statement of the Lemma hold for the arc $[e^-,e^+]$. By construction, $[e^-,e^+]$ contains all edges in B^- and B^+, thus the second property holds. By the definition of e^-, for every edge e' of $[e^-,e]$ we have $|P(e',\alpha)\cap P(e,\alpha)| \geq \left\lceil \frac{\mu(e,\alpha)}{2} \right\rceil$, and by definition of e^+, the same holds for every edge e' of $[e,e^+]$. Therefore, the second inequality of the first property is also satisfied. The first inequality is trivial since $\mu(e',\alpha)=|P(e',\alpha)|$. ■

A. Constant Price of Anarchy for Heavy Instances

We next 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 yields as well a constant 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 we proved in Lemma 8 in order to establish the existence of a heavily loaded edge in S-PMC(RING) games with $\hat{\mu}\geq w$.

Lemma 9. *In every* S-PMC(RING) *game* $\langle G,P,W\rangle$ *with* $\hat{\mu}\geq w$ *there is an edge with load at least* $\frac{\hat{\mu}\cdot w}{4}$.

Proof: Let $\mathbf{c}$ be a worst-case Nash equilibrium for game $\langle G,P,W\rangle$. We presently define a sequence of quintuples

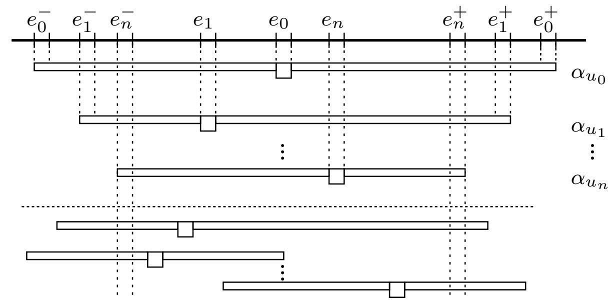

Fig. 4. The path structure implied in the proof of Lemma 9.

$\{(e_i,\alpha_{u_i},e_i^-,e_i^+,A_i)\}_{i\geq 0}$. Let e_0 and α_{u_0} be an edge and color, respectively, such that $\mu(e_0,\alpha_{u_0})=\hat{\mu}$. Let $[e_0^-,e_0^+]$ be an arc that satisfies the properties of Lemma 8 applied on edge e_0 and color α_{u_0} and let $A_0=\{\alpha_{u_0}\}$.

Now, for $i\geq 0$, given $\{(e_i,\alpha_{u_i},e_i^-,e_i^+,A_i)\}$, define e_{i+1} and $\alpha_{u_{i+1}}$ to be some edge in $[e_i^-,e_i^+]$ and color in $W\setminus A_i$, respectively, with the following properties:

1) $\mu(e_{i+1},\alpha_{u_{i+1}})\geq \mu(e_i,\alpha_{u_i})-1$ and
2) the application of Lemma 8 on edge e_{i+1} and color $\alpha_{u_{i+1}}$ yields an arc $[e_{i+1}^-,e_{i+1}^+]$ which is a subset of $[e_i^-,e_i^+]$.

Finally, define $A_{i+1}=A_i\cup\{\alpha_{u_{i+1}}\}$. This sequence of quintuples is defined up to $i=n\leq w-1$, at which point either $A_n=W$ or the application of Lemma 8 on any edge e_{n+1} in arc $[e_n^-,e_n^+]$ and any color $\alpha_{u_{n+1}}\in W\setminus A_n$ such that $\mu(e_{n+1},\alpha_{u_{n+1}})\geq\mu(e_n,\alpha_{u_n})-1$ fails to provide an arc which is contained in $[e_n^-,e_n^+]$. See Fig. 4 for an illustration of this structure.

From the definitions, it is immediately deduced that for all i,

$$\mu(e_i,\alpha_{u_i})\geq\hat{\mu}-i\ .$$

Therefore, by Lemma 8 for all edges e in $[e_i^-,e_i^+]$,

$$\mu(e,\alpha_{u_i})\geq\left\lceil\frac{\hat{\mu}-i}{2}\right\rceil\geq\frac{\hat{\mu}-i}{2}\ .$$

In particular, edges e_n^- and e_n^+ are included in all of the intervals $[e_i^-,e_i^+]$. Therefore, the load induced on edge e_n^- by paths colored with colors in A_n is:

$$\sum_{i=0}^{n}\mu(e_n^-,\alpha_{u_i})\geq\sum_{i=0}^{n}\frac{\hat{\mu}-i}{2}\geq\frac{(n+1)}{2}\cdot\hat{\mu}-\frac{n\cdot(n+1)}{4}\ . \tag{6}$$

Similarly for edge e_n^+:

$$\sum_{i=0}^{n}\mu(e_n^+,\alpha_{u_i})\geq\frac{(n+1)}{2}\cdot\hat{\mu}-\frac{n\cdot(n+1)}{4}\ .$$

Furthermore, the application of Lemma 8 on any color $\alpha_{u_{n+1}}\in W\setminus A_n$ and any edge e_{n+1} in $[e_n^-,e_n^+]$ such that $\mu(e_{n+1},\alpha_{u_{n+1}})\geq\mu(e_n,\alpha_{u_n})-1$ yields an arc $[e^-,e^+]$ not contained in $[e_n^-,e_n^+]$. This implies that at least half of the colors in $W\setminus A_n$ induce arcs that contain the same extremal edge of $[e_n^-,e_n^+]$ (let it be edge e_n^-, without loss of generality). Fig. 4 offers an illustration. Due to Lemma 8, for each such color α, each edge in $[e^-,e^+]$ – and therefore also edge e_n^- –

is used by at least $\left\lceil \frac{\mu(e_n,\alpha_{u_n})-1}{2} \right\rceil$ paths of color α. Since the number of colors in $W \setminus A_n$ is $w-(n+1)$, the load induced on edge e_n^- by paths colored with colors in $W \setminus A_n$ is:

$$\sum_{\alpha \in W\setminus A_n} \mu(e_n^-,\alpha) = \left\lceil \frac{|W \setminus A_n|}{2} \right\rceil \cdot \left\lceil \frac{\mu(e_n,\alpha_{u_n})-1}{2} \right\rceil \geq \frac{w-(n+1)}{2} \cdot \frac{\hat{\mu}-(n+1)}{2} \ . \quad (7)$$

By (6) and (7), the total load of edge e_n^- is:

$$L(e_n^-) = \sum_{i=0}^{n} \mu(e_n^-,\alpha_{u_i}) + \sum_{\alpha \in W\setminus A_n} \mu(e_n^-,\alpha) \geq \frac{\hat{\mu}\cdot w}{4} + \frac{n+1}{4} \cdot (\hat{\mu}-w+1) \ . \quad (8)$$

Since $\hat{\mu} \geq w$ and $n \geq 0$, we get from (8) that $L(e_n^-) \geq \frac{\hat{\mu}\cdot w}{4}$. ■

We can now prove a constant upper bound on the price of anarchy of games in S-PMC(RING: $L=\Omega(w^2)$).

Theorem 6. *The price of anarchy of any game in the class* S-PMC(RING) *is at most* $\max\{4, \frac{w^2}{L}\}$.

Proof: We distinguish between two cases: If $\hat{\mu} \geq w$, then by Lemma 9 we have a maximum load $L \geq \frac{\hat{\mu}\cdot w}{4}$. Therefore, the price of anarchy is bounded as follows:

$$\text{PoA} = \frac{\hat{\mu}}{\mu_{\text{OPT}}} \leq \frac{\hat{\mu}\cdot w}{L} \leq 4 \ ,$$

where for the first inequality we used Fact 1.

On the other hand, if $\hat{\mu} < w$, then we can bound the price of anarchy as follows:

$$\text{PoA} \leq \frac{\hat{\mu}\cdot w}{L} < \frac{w^2}{L} \ .$$

■

Corollary 10. *The price of anarchy of heavily loaded games defined on rings, i.e. of the class* S-PMC(RING: $L=\Omega(w^2)$), *is at most* $O(1)$. *In particular, the price of anarchy of games in* S-PMC(RING: $L \geq \frac{w^2}{4}$) *is at most* 4.

B. Unbounded Price of Anarchy for Light Instances

To complete our study of the price of anarchy in graphs of degree 2, we show that for arbitrarily small ε there exists an infinite subclass of S-PMC(CHAIN: $L=O(w^{2-\varepsilon})$) games whose price of anarchy is $\Omega(w^{\frac{\varepsilon}{2}})$. This implies that the price of anarchy can get arbitrarily large when the number of available colors increases, therefore the price of anarchy is unbounded for the classes S-PMC(CHAIN: $L=o(w^2)$) and, consequently, S-PMC(RING: $L=o(w^2)$).

Theorem 7. *For any fixed* ε *in the range* $0<\varepsilon\leq 1$, *there exists an infinite family of games in* S-PMC(CHAIN: $L=O(w^{2-\varepsilon})$) *with price of anarchy* $\Omega(w^{\frac{\varepsilon}{2}})$.

Proof: Fix some ε in the range $0<\varepsilon\leq 1$. We will construct a family $\{\mathcal{B}_\rho\}_{\rho\geq 4}$ of S-PMC(CHAIN: $L=O(w^{2-\varepsilon})$) games and a strategy profile for each game in the family, which

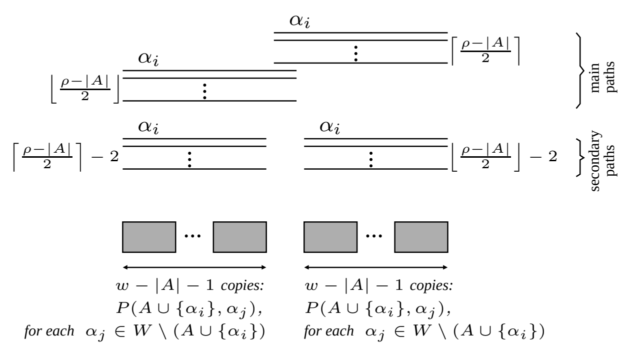

Fig. 5. The path set $P(A,\alpha_i)$, for $\alpha_i \in W\setminus A$, used in the construction of Theorem 7.

we will prove to be a Nash equilibrium, with the following properties:

1) the number of available colors is $w=\left\lceil \rho^{1+\frac{\varepsilon}{2-\varepsilon}} \right\rceil$,
2) the maximum load is $L=\frac{\rho^2-3\cdot\rho+12}{2}$, and
3) the maximum multiplicity of any color is $\mu_{\max}=\rho$.

First of all, observe that $L=O(\rho^2)=O(w^{2-\varepsilon})$. Therefore, the family indeed belongs to S-PMC(CHAIN: $L=O(w^{2-\varepsilon})$). Moreover, since the class S-PMC(CHAIN) is a subclass of S-PMC(ROOTED-TREE), by Theorem 2 the optimal strategy profile for game $\mathcal{B}_\rho$ has social cost $\mu_{\text{OPT}}=\left\lceil \frac{L}{w} \right\rceil < \frac{L}{w}+1$. Additionally, the cost of the worst-case Nash equilibrium must be $\hat{\mu}\geq\mu_{\max}$. The price of anarchy of game $\mathcal{B}_\rho$ is therefore:

$$\text{PoA}(\mathcal{B}_\rho) = \frac{\hat{\mu}}{\mu_{\text{OPT}}} \geq \frac{\mu_{\max}}{\frac{L}{w}+1} = \frac{w\cdot\mu_{\max}}{L+w} \geq \frac{\rho^{1+\frac{\varepsilon}{2-\varepsilon}}\cdot\rho}{\frac{\rho^2-3\cdot\rho+12}{2}+\rho^{1+\frac{\varepsilon}{2-\varepsilon}}+1} \ ,$$

where in the last step we used the fact that $\rho^{1+\frac{\varepsilon}{2-\varepsilon}} \leq w < \rho^{1+\frac{\varepsilon}{2-\varepsilon}}+1$. Because $\varepsilon\leq 1$, we get that $\text{PoA}(\mathcal{B}_\rho)=\Omega(\rho^{\frac{\varepsilon}{2-\varepsilon}})=\Omega(w^{\frac{\varepsilon}{2}})$.

Construction of $\mathcal{B}_\rho$*:* Given the parameters ε and ρ, we describe the construction of an S-PMC(CHAIN) game using the path set $P(A,\alpha)$ illustrated in Fig. 5 as a building block. We now describe the structure of $P(A,\alpha)$, along with a coloring of the paths it contains.

Let W be a set of available colors of size $\left\lceil \rho^{1+\frac{\varepsilon}{2-\varepsilon}} \right\rceil$. We will define $P(A,\alpha)$ for $|A|\leq\rho-3$, $A\subseteq W$, and $\alpha\in W\setminus A$. Whenever $|A|\leq\rho-4$, $P(A,\alpha)$ is recursively defined to be a path set consisting of:

- The *main paths*: these are the $\rho-|A|$ paths of color α, arranged as shown in the top part of Fig. 5. They all share a common edge, henceforth called the *central edge* for this copy of $P(A,\alpha)$. Half of them extend to the left of the central edge, and the rest extend to the right.
- The *secondary paths*: these are the $\rho-|A|-4$ paths of color α, arranged as illustrated in the central part of Fig. 5. These paths do not use the central edge and half of them extend to the left of the central edge, and the rest extend to the right.
- Two copies of $P(A\cup\{\alpha\},\alpha')$ for every $\alpha'\in W\setminus(A\cup\{\alpha\})$, one of them on each side of the central edge, as per the bottom part of Fig. 5.

The base case of this recursive construction occurs when $|A| = \rho - 3$. In this case, the path set $P(A, \alpha)$ is defined on a chain of length $\ell = (w - \rho + 2) \cdot (w - \rho + 1)$. Let $e_0, e_1, \ldots, e_{\ell-1}$ be the edges of this chain, in order from left to right. The path set $P(A, \alpha)$ contains:

- three paths of color α spanning edges e_0 up to $e_{\ell-1}$,
- for each color $\beta_i \in W \setminus (A \cup \{\alpha\})$, two paths of color β_i and length $(w-\rho+1)$ spanning edges $e_{i\cdot(w-\rho+1)}$ up to $e_{(i+1)\cdot(w-\rho+1)-1}$, where $\beta_0, \ldots, \beta_{w-\rho+1}$ is an arbitrary enumeration of the colors in $W \setminus (A \cup \{\alpha\})$, and
- for each color β_i in the previous enumeration, for each color $\gamma_{i,j} \in W \setminus (A \cup \{\alpha, \beta_i\})$, one path of color $\gamma_{i,j}$ and length one, defined on edge $e_{i\cdot(w-\rho+1)+j}$, where $\gamma_{i,0}, \ldots, \gamma_{i,w-\rho}$ is an arbitrary enumeration of the colors in $W \setminus (A \cup \{\alpha, \beta_i\})$.

We now claim that the S-PMC(CHAIN) game $\langle G, P(\emptyset, \alpha_1), W\rangle$, where G is a chain long enough to accommodate all paths of $P(\emptyset, \alpha_1)$, is a game in S-PMC(CHAIN: $L = O(w^{2-\varepsilon})$) with the desired properties. First observe that the maximum color multiplicity $\mu_{\max}$ occurs, by construction, on the central edge of the chain. This edge is used by exactly ρ paths of color α_1, thus $\mu_{\max} = \rho$. The number of available colors has been fixed to exactly $w = \lceil \rho^{1+\frac{\varepsilon}{2-\varepsilon}} \rceil$. Finally, regarding the maximum load, observe that the main and secondary paths of each copy of $P(A, \alpha)$ induce a load of $\rho - |A| - 2$, whereas the base case of the construction induces a constant load of 6. Therefore, the maximum load is:

$$L = 6 + \sum_{i=0}^{\rho-4} (\rho - i - 2) = \frac{\rho^2 - 3 \cdot \rho + 12}{2} \ .$$

It remains to be shown that the coloring described above for path set $P(\emptyset, \alpha_1)$ is indeed a Nash equilibrium. Consider a particular copy of $P(A, \alpha)$ with $|A| \leq \rho - 4$, included somewhere in the path set. We remark immediately that, by construction, the main and secondary paths do not overlap with any other path of color α in the path set. Therefore, the main paths have a disutility of $\rho - |A|$ and the secondary paths have a disutility of $\rho - |A| - 2$.

Now, for all colors $\alpha' \in W \setminus (A \cup \{\alpha\})$, any main or secondary path of $P(A, \alpha)$ overlaps with a copy of $P(A \cup \{\alpha\}, \alpha')$, thus it contains some edge on which the multiplicity of color α' is $\rho - (|A| + 1) = \rho - |A| - 1$. This proves that all main and secondary paths in $P(A, \alpha)$ are blocked from switching to any color $\alpha' \in W \setminus (A \cup \{\alpha\})$.

Regarding colors $\alpha' \in A$, observe that the copy of $P(A, \alpha)$ under consideration is itself included in a sequence of bigger copies $P(A_1, \alpha'_1), P(A_2, \alpha'_2), \ldots, P(A_x, \alpha'_x)$, where $A_1 = A \setminus \{\alpha'_1\}$ and $A_{i+1} = A_i \setminus \{\alpha'_{i+1}\}$, ending at $A_x = \emptyset$, $\alpha'_x \equiv \alpha_1$. Therefore, for any color $\alpha'_i \in A$, any edge of any main or secondary path of $P(A, \alpha)$ has a multiplicity of color α'_i equal to $\rho - (|A| - i) - 2 = \rho - |A| + i - 2 \geq \rho - |A| - 1$. Therefore, all main and secondary paths in $P(A, \alpha)$ are also blocked from switching to any color $\alpha' \in A$.

Finally, it is straightforward to verify that all paths contained in the base case of the construction contain the required blocking edges, and thus the game is in Nash equilibrium. ∎

VII. CONCLUSION

In this work we have proposed a framework for studying non-cooperative wavelength assignment in multifiber optical networks, namely SELFISH PATH MULTICOLORING games. The results we obtained in Section VI suggest an efficient decentralized protocol for wavelength assignment in ring networks where wavelengths are scarce – a safe assumption in practically relevant scenarios: Any Nash equilibrium reached by the players is guaranteed to have a cost of at most 4 times that of the cost of an optimal solution.

On the other hand, the situation in more general networks is much less gratifying: In Section V we show that the price of anarchy can grow unbounded even in tree networks of maximum degree 3. Evidently, the player-charging mechanism of this model is inadequate to steer players towards Nash equilibria of low fiber cost in such networks. It is therefore a promising direction for further research to investigate the quality of equilibria obtained under different player cost functions.

ACKNOWLEDGMENT

The authors wish to thank the anonymous reviewers for their thorough comments which helped to significantly improve the presentation of the article, as well as clarify a number of technical points.

REFERENCES

[1] E. Altman, T. Boulogne, R. E. Azouzi, T. Jiménez, and L. Wynter, "A survey on networking games in telecommunications," *Computers & OR*, vol. 33, pp. 286–311, 2006.
[2] J. Huang, D. Palomar, N. Mandayam, J. Walrand, S. Wicker, and T. Başar, Eds., *the IEEE JSAC September 2008 special issue on Game Theory in Communication Systems*. IEEE, 2008.
[3] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, Eds., *Algorithmic Game Theory*. Cambridge University Press, 2007.
[4] V. Bilò and L. Moscardelli, "The price of anarchy in all-optical networks," in *SIROCCO*, ser. Lect. Notes Comput. Sci., R. Kralovic and O. Sýkora, Eds., vol. 3104. Springer, 2004, pp. 13–22.
[5] V. Bilò, M. Flammini, and L. Moscardelli, "On Nash equilibria in non-cooperative all-optical networks," in *STACS*, ser. Lect. Notes Comput. Sci., V. Diekert and B. Durand, Eds., vol. 3404. Springer, 2005, pp. 448–459.
[6] G. F. Georgakopoulos, D. J. Kavvadias, and L. G. Sioutis, "Nash equilibria in all-optical networks," *Discrete Math.*, vol. 309, no. 13, pp. 4332–4342, 2009.
[7] I. Milis, A. Pagourtzis, and K. Potika, "Selfish routing and path coloring in all-optical networks," in *CAAN*, ser. Lect. Notes Comput. Sci., J. C. M. Janssen and P. Pralat, Eds., vol. 4852. Springer, 2007, pp. 71–84.
[8] E. Koutsoupias and C. H. Papadimitriou, "Worst-case equilibria," in *STACS*, ser. Lect. Notes Comput. Sci., C. Meinel and S. Tison, Eds., vol. 1563. Springer, 1999, pp. 404–413.
[9] ——, "Worst-case equilibria," *Computer Science Review*, vol. 3, no. 2, pp. 65–69, 2009.
[10] C. Nomikos, A. Pagourtzis, and S. Zachos, "Routing and path multicoloring," *Inform. Process. Lett.*, vol. 80, no. 5, pp. 249–256, 2001.
[11] J. Nash, "Non-cooperative games," *Ann. Math.*, vol. 54, no. 2, pp. 286–295, 1951.
[12] A. Orda, R. Rom, and N. Shimkin, "Competitive routing in multiuser communication networks," *IEEE/ACM Trans. Netw.*, vol. 1, no. 5, pp. 510–521, 1993.
[13] M. Andrews and L. Zhang, "Minimizing maximum fiber requirement in optical networks," *J. Comput. Syst. Sci.*, vol. 72, no. 1, pp. 118–131, 2006.

[14] L. Margara and J. Simon, "Wavelength assignment problem on all-optical networks with k fibres per link," in *ICALP*, ser. Lect. Notes Comput. Sci., U. Montanari, J. D. P. Rolim, and E. Welzl, Eds., vol. 1853. Springer, 2000, pp. 768–779.
[15] G. Li and R. Simha, "On the wavelength assignment problem in multifiber WDM star and ring networks," *IEEE/ACM Trans. Netw.*, vol. 9, no. 1, pp. 60–68, 2001.
[16] A. K. Somani, M. Mina, and L. Li, "On trading wavelengths with fibers: a cost-performance based study," *IEEE/ACM Trans. Netw.*, vol. 12, no. 5, pp. 944–951, 2004.
[17] L. Li and A. K. Somani, "Blocking performance of fixed-paths least-congestion routing in multifibre WDM networks," *Int. J. Communication Systems*, vol. 15, no. 2-3, pp. 143–159, 2002.
[18] G. Christodoulou and E. Koutsoupias, "The price of anarchy of finite congestion games," in *STOC*, H. N. Gabow and R. Fagin, Eds. ACM, 2005, pp. 67–73.
[19] ——, "On the price of anarchy and stability of correlated equilibria of linear congestion games," in *ESA*, ser. Lecture Notes in Computer Science, G. S. Brodal and S. Leonardi, Eds., vol. 3669. Springer, 2005, pp. 59–70.
[20] C. Busch and M. Magdon-Ismail, "Atomic routing games on maximum congestion," *Theor. Comput. Sci.*, vol. 410, no. 36, pp. 3337–3347, 2009.
[21] R. Banner and A. Orda, "Bottleneck routing games in communication networks," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 6, pp. 1173–1179, 2007.
[22] E. Anshelevich, A. Dasgupta, J. M. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden, "The price of stability for network design with fair cost allocation," *SIAM J. Comput.*, vol. 38, no. 4, pp. 1602–1623, 2008.
[23] M. Andrews and L. Zhang, "Complexity of wavelength assignment in optical network optimization," *IEEE/ACM Trans. Netw.*, vol. 17, no. 2, pp. 646–657, 2009.
[24] ——, "Wavelength assignment in optical networks with fixed fiber capacity," in *ICALP*, ser. Lect. Notes Comput. Sci., J. Díaz, J. Karhumäki, A. Lepistö, and D. Sannella, Eds., vol. 3142. Springer, 2004, pp. 134–145.
[25] T. Erlebach, A. Pagourtzis, K. Potika, and S. Stefanakos, "Resource allocation problems in multifiber WDM tree networks," in *WG*, ser. Lect. Notes Comput. Sci., H. L. Bodlaender, Ed., vol. 2880. Springer, 2003, pp. 218–229.
[26] C. Nomikos, A. Pagourtzis, K. Potika, and S. Zachos, "Routing and wavelength assignment in multifiber wdm networks with non-uniform fiber cost," *Computer Networks*, vol. 50, no. 1, pp. 1–14, 2006.
[27] M. Andrews and L. Zhang, "Bounds on fiber minimization in optical networks with fixed fiber capacity," in *INFOCOM*. IEEE, 2005, pp. 409–419.
[28] A. Pagourtzis, K. Potika, and S. Zachos, "Path multicoloring with fewer colors in spiders and caterpillars," *Computing*, vol. 80, no. 3, pp. 255–274, 2007.
[29] T. Harks, M. Hoefer, M. Klimm, and A. Skopalik, "Computing pure Nash and strong equilibria in bottleneck congestion games," in *ESA (2)*, ser. Lecture Notes in Computer Science, M. de Berg and U. Meyer, Eds., vol. 6347. Springer, 2010, pp. 29–38.
[30] H. Ackermann, H. Röglin, and B. Vöcking, "On the impact of combinatorial structure on congestion games," *J. ACM*, vol. 55, no. 6, 2008.
[31] S. Chien and A. Sinclair, "Convergence to approximate Nash equilibria in congestion games," in *SODA*, N. Bansal, K. Pruhs, and C. Stein, Eds. SIAM, 2007, pp. 169–178.
[32] C. L. Monma and V. K.-W. Wei, "Intersection graphs of paths in a tree," *J. Comb. Theory B*, vol. 41, no. 2, pp. 141–181, 1986.

Evangelos Bampas received the Dipl.Ing. and Ph.D. degree in electrical & computer engineering from the National Technical University of Athens (NTUA), Greece, in 2004 and 2009, respectively.

He is currently a Research Associate at the Computation & Reasoning Laboratory of NTUA. During the Spring of 2009 he was a Research Intern at INRIA Bordeaux, France (team CEPAGE), where he continued during the 2009–2010 academic year as a Postdoctoral Fellow. His research interests lie in the areas of optical networks, mobile agent computing, computational complexity theory, and algorithmic game theory.

Dr. Bampas is a member of the Technical Chamber of Greece and the European Association for Theoretical Computer Science.

Aris Pagourtzis holds a Dipl.Ing. degree in electrical engineering (1989) and a Dr.Eng. degree in electrical and computer engineering (1999), both from the National Technical University of Athens.

He is currently an Assistant Professor at the School of Electrical and Computer Engineering of NTUA, and has previously held academic posts at the University of Ioannina, the University of Liverpool, ETH Zürich, and the University of Athens. He has served in the program committee of several computer science conferences. His main research interests lie in the areas of algorithms for optical networks, distributed protocols for ad hoc radio networks, approximability, computational complexity of counting problems, and algorithmic game theory.

Dr. Pagourtzis is a member of EATCS and the Technical Chamber of Greece, and has served as a reviewer for IEEE/ACM Transactions on Networking and IEEE Transactions on Computers.

George Pierrakos received his Dipl.Ing. degree in electrical and computer engineering from the National Technical University of Athens in 2007. He then received a M.S. degree in mathematics from the University of Athens in 2009. He is currently a Ph.D. candidate in computer science at the department of electrical engineering and computer sciences at UC Berkeley.

He has worked as a Research Assistant for the Computation and Reasoning Laboratory at the National Technical University of Athens for 1 year, and has since worked as a Research or Teaching Assistant for the department of electrical engineering and computer sciences at UC Berkeley. He is currently an Intern with Microsoft Research New England, in Cambridge, MA. His research interests lie in the area of theoretical computer science and algorithmic game theory.

Mr. Pierrakos is the recipient of a Bodossaki Foundation Scholarship, an Onassis Public Benefit Foundation Scholarship and a UC Regents Fellowship.

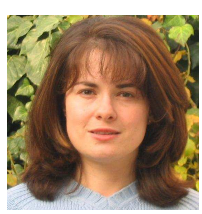

Katerina Potika received the Dipl.Ing. degree in electrical and computer engineering and the Dr.Eng. degree in computer science from the National Technical University of Athens (NTUA), Athens, Greece, in 1997 and 2004, respectively.

She is currently a Research Associate at NTUA and a Lecturer at the San Jose State Univesity. She has previously been a lecturer at the University of California (Santa Cruz), the University of Athens, the Athens University of Economics and Business, and a Visiting Researcher at the University of Konstanz and at the Swiss Federal Institute of Technology Zürich (ETH Zürich). Her research interests include graph algorithms, graph drawing, optical networks and game theory.

Dr. Potika is a member of the Technical Chamber of Greece.