

On Nash Equilibria in Non-Cooperative All-Optical Networks [†]

Vittorio Bilò ¹, Michele Flammini ² and Luca Moscardelli ^{3,*}

¹ Dipartimento di Matematica e Fisica “Ennio De Giorgi”, University of Salento, Provinciale Lecce-Arnesano, P.O. Box 193, 73100 Lecce, Italy; vittorio.bilo@unisalento.it

² GSSI Institute, Viale Francesco Crispi 7, 67100 L'Aquila, Italy; michele.flammini@gssi.it

³ Dipartimento di Economia, University of Chieti-Pescara, Viale Pindaro 42, 65125 Pescara, Italy

* Correspondence: luca.moscardelli@unich.it

[†] Extended abstracts of this work appeared in the Proceedings of the 11th Colloquium on Structural Information and Communication Complexity (SIROCCO 2004), and in the Proceedings of the 22nd International Symposium on Theoretical Aspects of Computers (STACS 2005).

Abstract: We consider the problem of determining a routing in all-optical networks, in which some couples of nodes want to communicate. In particular, we study this problem from the point of view of a network provider that has to design suitable payment functions for non-cooperative agents, corresponding to the couples of nodes wishing to communicate. The network provider aims at inducing stable routings (i.e., routings corresponding to Nash equilibria) using a low number of wavelengths. We consider three different kinds of local knowledge that agents may exploit to compute their payments, leading to three corresponding information levels. Under complete information, the network provider can design a payment function, inducing the agents to reach a Nash equilibrium mirroring any desired routing. If the price to an agent is computed only as a function of the wavelengths used along connecting paths (minimal level) or edges (intermediate level), the most reasonable functions either do not admit Nash equilibria or admit very inefficient ones, i.e., with the largest possible price of anarchy. However, by suitably restricting the network topology, a constant price of anarchy for chains and rings and a logarithmic one for trees can be obtained under the minimal and intermediate levels, respectively.



Citation: Bilò, V.; Flammini, M.; Moscardelli, L. On Nash Equilibria in Non-Cooperative All-Optical Networks. *Algorithms* **2021**, *14*, 15. <https://doi.org/10.3390/a14010015>

Received: 10 December 2020

Accepted: 6 January 2021

Published: 9 January 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Keywords: optical networks; WDM; pricing of optical network services; selfish routing; nash equilibria; price of anarchy; algorithmic game theory

1. Introduction

Due to the possibility of managing thousands of users, covering wide areas and providing a bandwidth which is orders of magnitude faster than traditional networks, all-optical networks are widely exploited nowadays, and are considered to be used also in the future of communication networks. In fact, wavelength division multiplexing (WDM) [1] allows one to exploit the high bandwidth provided by all-optical networks by partitioning it a large number of parallel high speed channels along the same optical fiber (see [2,3] for a survey of the main related results). Motivated by this fact, in this paper, we investigate the existence and performance of Nash equilibria in all-optical networks populated by non-cooperative selfish agents corresponding to point-to-point communication requests. Every agent is interested only in the minimization of its own cost, and, given a graph modelling the optical network, has to choose among a set of possible paths (also called routing strategy) connecting its endpoints. In particular, in our model, a service provider has to design suitable payment functions, charging each agent a cost for her chosen routing strategy. We consider three different kinds of local knowledge that agents may exploit to compute their payments, leading to three corresponding information levels that will be detailed in the following. It is worth noticing that any request is willing to be rerouted each time it may be served by a cheaper path: the evolution of the network can be modelled by the dynamics of a multi-agent system or multi-player game. A routing solution, that

is an assignment of paths and colors (a color represents a wavelength) to the requests, in which no request can lower its cost by choosing a different strategy, is said to be a pure Nash equilibrium.

1.1. Related Work

Several routing games have been shown to possess pure Nash equilibria or to converge to a pure Nash equilibrium independently from their starting state (see [4–14] for some pioneering work and [15] for an introduction to algorithmic game theory). The price of anarchy has been introduced in [12] for the first time in the literature; it measures the loss of the global performance of the system at equilibrium due to the lack of cooperation among the agents, and is defined as the worst-case ratio between the total cost of a Nash equilibrium and that of an optimal solution. One of the most studied and interesting research areas lying on the boundary between computer science and game theory consists of bounding the price of anarchy of selfish routing under different models; see, for example [13,14].

Game theory has been applied to optical networks in various contexts (e.g., see [16–18]). For instance, it is worth noting that the optimization of the network cost due to expensive hardware components (e.g., add-drop multiplexer) has been widely investigated in the literature (see [19–22]), also considering the effect of collusion among agents (see [23]). Moreover, in [16], the authors deal with a setting very similar to the one addressed in this paper; they consider several natural cost functions, and for each of them, they examine the problem of the existence of pure Nash equilibria, and of the complexity of recognizing and computing such equilibria.

1.2. Our Contribution

As a first contribution, we determine three notable information levels for the computation of the payment functions that have to be designed by the network provider. In particular, each agent has knowledge of all the other agents and their routing strategies under the complete information level of the wavelengths used along any given edge under the intermediate level, and finally, only of the wavelengths used along any given path connecting its source and destination under the minimal level. Notice that even the minimal level contains the information needed in order to select an available wavelength, without interfering with other requests.

We first show that, under complete level, the network provider can design a payment function inducing the agents to reach a Nash equilibrium mirroring any desired routing. This allows the use of a centralized algorithm for determining an (almost) efficient routing assignment which can be then efficiently enforced to the agents. We then turn our attention to the remaining two levels of information: unfortunately, it holds that the most reasonable functions either do not admit equilibria or induce games with a very high (actually the worst possible) price of anarchy; that is, the resulting Nash equilibrium could assign a different color to each agent. Given this negative result, we focus on more specific network topologies. In particular, we show that a price of anarchy of 8 and 16 can be obtained for chains and rings, respectively, under the minimal level of information. These bounds can be further reduced to $3 + \epsilon$ and $6 + \epsilon$ under the intermediate level, with ϵ converging to 0 as the load increases. Finally, a price of anarchy logarithmic in the number of agents is shown to hold for tree networks under the minimal level of information.

1.3. Paper Organization

The paper is organized as follows. Section 2 provides the basic definitions and notation. In Section 3 we give some preliminary results for our model: in particular, we present the results related to the complete information level and we also define some suitable classes of payment functions able to characterize convergent games in the context of more restrictive information levels. Sections 4 and 5 are devoted to the study of payment functions for the minimal and intermediate levels of information: in particular, in Section 4, results

relative to general networks are presented, while in Section 5, results for specific network topologies, that is chains, rings and trees, are provided. Finally, Section 6 gives some conclusive remarks and discusses some open questions.

2. The Model

An all-optical network is modelled by an undirected graph $G = (V, E)$, where nodes in V represent sites and undirected edges in E bidirectional optical fiber links between the sites. Given any two nodes $x, y \in V$, a communication request between x and y is denoted as $\{x, y\}$. A communication instance in G is a multiset of requests I (i.e., I can contain multiple requests between the same pair of nodes). A path system P for an instance I in G is a set of paths containing a distinguished simple connecting path in G for each request in I . A solution $\mathcal{R}(G, I)$ for an instance I in G , \mathcal{R} for short, is a pair $(P_{\mathcal{R}}, c_{\mathcal{R}})$ in which $P_{\mathcal{R}}$ is a path system for I and $c_{\mathcal{R}} : I \rightarrow W$ (with $W = \mathbb{N}^+$ being the set of wavelengths) is a function associating a wavelength or color to each request in I . Let $p_{\mathcal{R}}(\{x, y\}) \in P_{\mathcal{R}}$ denote the path connecting $\{x, y\}$ in \mathcal{R} and $|p_{\mathcal{R}}(\{x, y\})|$ be its length in terms of number of edges. A solution \mathcal{R} is feasible or is a routing for I in G if $c_{\mathcal{R}}(\{x_1, y_1\}) \neq c_{\mathcal{R}}(\{x_2, y_2\})$ for any two requests $\{x_1, y_1\} \in I$ and $\{x_2, y_2\} \in I$ whose connecting paths $p_{\mathcal{R}}(\{x_1, y_1\})$ and $p_{\mathcal{R}}(\{x_2, y_2\})$ share an edge in G , i.e., the associated colors are different.

Let $\omega_{\mathcal{R}}(G, I)$ be the number of colors used by the routing \mathcal{R} for I in G and $\omega(G, I) = \min_{\mathcal{R}} \omega_{\mathcal{R}}(G, I)$ be the minimum number of colors used by any routing for I . Analogously, let $\pi_{\mathcal{R}}(G, I) = \max_{e \in E} |\{p \in P_{\mathcal{R}} | e \in p\}|$ be the maximum load of an edge in G induced by the routing \mathcal{R} and $\pi(G, I) = \min_{\mathcal{R}} \pi_{\mathcal{R}}(G, I)$ be the minimum maximum load of the routings for I in G , also called load of I in G . Clearly, since all the requests sharing an edge must have different colors, $\omega_{\mathcal{R}}(G, I) \geq \pi_{\mathcal{R}}(G, I)$ for every \mathcal{R} and thus $\omega(G, I) \geq \pi(G, I)$. Given that the optical spectrum is a scarce resource, in this paper, we aim at determining routings using a number of wavelengths close to the minimum one, i.e., $\omega(G, I)$.

In order to model our non-cooperative environment, we assume that each communication request $\{x, y\} \in I$ is issued and handled by an agent α ; for the sake of simplicity, we consider it as coincident with the request, that is $\alpha = \{x, y\}$. Our goal is to assign payments (or costs) to the agents so that they are induced to choose in a non-cooperative way a good solution, i.e., a routing with a small number of colors. A payment function *price* associates each possible routing and agent to a cost; in particular, $price_{\mathcal{R}}(\alpha)$ is the cost that agent $\alpha \in I$ has to pay to the network provider in order to obtain the asked service if the routing \mathcal{R} is adopted. In order to give our payment functions a higher generality, we assume the existence of a non-decreasing function $f : W \rightarrow \mathbb{R}^+$ associating a cost to every color. Notice that f is non-decreasing because, in such a way, it is possible to implicitly model the increasing cost incurred by the network provider when it has implemented a routing using up to a given wavelength.

A game $\mathcal{G} = (G, I, price)$ among the $|I|$ agents belonging to I on the network G induced by the payment function *price* is defined as the set of agents' strategies $P^{\alpha} \times W$, where P^{α} is the set of connecting paths for α and the utility function $u(\alpha) = -price(\alpha)$. A routing \mathcal{R} is a (pure) Nash equilibrium if, and only if, for any agent α and routing \mathcal{R}' differing from \mathcal{R} only for the path and/or the color associated to α , it holds that $price_{\mathcal{R}}(\alpha) \leq price_{\mathcal{R}'}(\alpha)$. Denoted as \mathcal{N} , the set of the routings at Nash equilibrium, the price of anarchy of the game \mathcal{G} is defined as the worst case ratio $\rho = \sup_{\mathcal{R} \in \mathcal{N}} \frac{\omega_{\mathcal{R}}(G, I)}{\omega(G, I)}$ among all the possible games \mathcal{G} .

The evolution of the game \mathcal{G} is a sequence of moves $(\alpha, p_{old}, color_{old}, p_{new}, color_{new})$ starting from an arbitrary routing, where α is an agent and $(p_{old}, color_{old}) \in P^{\alpha} \times W$ and $(p_{new}, color_{new}) \in P^{\alpha} \times W$ are the old and the new strategies of α , respectively. The Nash dynamics of the game \mathcal{G} are the directed graph (Φ, M) , where Φ is the set of possible routings for G and I and there exists an arc $(\mathcal{R}_1, \mathcal{R}_2) \in M$ if there exists a selfish move from \mathcal{R}_1 to \mathcal{R}_2 , i.e., there exists an agent α such that $price_{\mathcal{R}_2}(\alpha) < price_{\mathcal{R}_1}(\alpha)$. If the Nash dynamics are acyclic for every G and I , the game \mathcal{G} is said to be convergent. Analogously, any payment function *price* inducing only convergent games is said to be convergent.

We say that game \mathcal{G} converges to routing \mathcal{R} whenever \mathcal{G} is convergent and admits the unique pure Nash equilibrium \mathcal{R} . It is worth noting that, even if the Nash dynamics are cyclic, the game might admit a Nash equilibrium.

Table 1 summarizes all the notations defined above. For ease of notation, when clear from the context, in the following, we will drop the index \mathcal{R} from the notation, thus, for instance, writing $price(\alpha)$, $p(\alpha)$, $c(\alpha)$, $\sigma(e)$ and $\bar{\sigma}(p)$ in place of $price_{\mathcal{R}}(\alpha)$, $p_{\mathcal{R}}(\alpha)$, $c_{\mathcal{R}}(\alpha)$, $\sigma_{\mathcal{R}}(e)$ and $\bar{\sigma}_{\mathcal{R}}(p)$. Before concluding the section, we finally observe that the payment function $price_{\mathcal{R}}$, in a strongly distributed non-cooperative environment, has to be computed by each single agent requiring the communication service. However, the level of global information they have can be limited by technological constraints, as well as privacy policies carried out by the service provider or simply enforced by the law. Therefore, in general, $price_{\mathcal{R}}$ is not computed, starting from the instance I and the routing \mathcal{R} , but on a more restricted set of information induced by them. To this aim, we introduce the concepts of states and levels of information. The edge state (resp. path state) of the network G induced by a routing \mathcal{R} for I is a function $\sigma_{\mathcal{R}} : E \rightarrow 2^W$ (resp. $\bar{\sigma}_{\mathcal{R}} : \mathcal{P} \rightarrow 2^W$, where \mathcal{P} is the set of all the simple paths in G) associating to every edge $e \in E$ (resp. path $p \in \mathcal{P}$) the set of the wavelengths used along e (resp. p). It is then possible to distinguish among three basic levels of information:

- Minimal. Each agent $\alpha = \{x, y\}$ knows the available wavelengths along any given path connecting x to y , that is, the function $price_{\mathcal{R}}$ can be computed, even knowing only the restriction of the path state $\bar{\sigma}_{\mathcal{R}}(\alpha)$ on the set of the paths from x to y .
- Intermediate. Each agent information knows the wavelengths available along any given edge, that is, the function $price_{\mathcal{R}}$ can be computed, even knowing only the edge state $\sigma_{\mathcal{R}}$.
- Complete. Each agent $\alpha = \{x, y\}$ knows the instance I and the routing \mathcal{R} , that is, the function $price$ is not restricted.

Table 1. Notation used in this paper.

$G = (V, E)$	\triangleq	The graph modelling the optical network, with V being the set of sites and E , the set of bidirectional optical links between the sites.
\mathcal{P}	\triangleq	The set of all simple paths in G .
$\{x, y\}$	\triangleq	A communication request between sites x and y ($x, y \in V$).
I	\triangleq	A multiset of requests defining a communication instance.
P	\triangleq	A path system.
W	\triangleq	The set of available wavelengths or colors.
$\mathcal{R}(G, I) = (P_{\mathcal{R}}, c_{\mathcal{R}})$	\triangleq	$\mathcal{R}(G, I)$ is a routing (i.e., a solution) for an instance I in G ; $P_{\mathcal{R}}$ is a path system for I ; $c_{\mathcal{R}}$ is a function associating a color to each request.
$p_{\mathcal{R}}(\{x, y\})$	\triangleq	The path connecting $\{x, y\}$ in \mathcal{R} .
$\omega_{\mathcal{R}}(G, I)$	\triangleq	The number of colors used by the routing \mathcal{R} for I in G .
$\omega(G, I)$	\triangleq	The minimum number of colors used by any routing for I in G .
$\pi_{\mathcal{R}}(G, I)$	\triangleq	The maximum load of an edge in G induced by the routing \mathcal{R} for instance I .
$\pi(G, I)$	\triangleq	The minimum maximum load of the routings for I in G , also called load of I in G .
$\alpha = \{x, y\}$	\triangleq	A generic agent associated to request $\{x, y\}$.
$price_{\mathcal{R}}(\alpha)$	\triangleq	The cost that agent $\alpha \in I$ has to pay to the network provider in order to obtain the asked service if the routing \mathcal{R} is adopted.
$f : W \rightarrow \mathbb{R}^+$	\triangleq	A function associating a cost to every color.
$\mathcal{G} = (G, I, price)$	\triangleq	A game among agents in I on the network G .
P^α	\triangleq	The set of connecting paths for agent α .
$u(\alpha)$	\triangleq	The utility of agent α .
\mathcal{N}	\triangleq	The set of Nash equilibria of a game.

Clearly, even if not explicitly mentioned, in any level of information, for each agent α , $price_{\mathcal{R}}(\alpha)$ can depend also on the wavelength $c(\alpha)$ and the path $p(\alpha)$ assigned to α in \mathcal{R} . Recall that the players move sequentially and can choose, as their current strategy, only valid strategies with respect to the current routing. Notice also that even with minimal information, any agent is always able to compute a valid wavelength for the chosen path, so as to not interfere with any other agent, and thus not compromising the feasibility of the solution.

3. Preliminaries

In the first part of this section, we prove that, under complete information, any centralized algorithm \mathcal{A} for the all-optical routing problem can be suitably used in a non-cooperative environment for achieving a price of anarchy either equal to 1 or k , according to whether \mathcal{A} is optimal or k -approximating, respectively.

Theorem 1. *Given any routing $\tilde{\mathcal{R}}$ for a game \mathcal{G} , there exists a payment function letting \mathcal{G} converge to $\tilde{\mathcal{R}}$ in at most $3|I|$ steps.*

Proof. Given any routing \mathcal{R} , the payment function is defined as follows.

$$price_{\mathcal{R}}(\alpha) = \begin{cases} 0 & \text{if } p_{\mathcal{R}}(\alpha) = p_{\tilde{\mathcal{R}}}(\alpha) \wedge c_{\mathcal{R}}(\alpha) = c_{\tilde{\mathcal{R}}}(\alpha), \\ 2 & \text{if } (p_{\mathcal{R}}(\alpha) \neq p_{\tilde{\mathcal{R}}}(\alpha) \vee c_{\mathcal{R}}(\alpha) \neq c_{\tilde{\mathcal{R}}}(\alpha)) \wedge c_{\mathcal{R}}(\alpha) \leq \omega_{\tilde{\mathcal{R}}}(G, I), \\ 1 & \text{if } (p_{\mathcal{R}}(\alpha) \neq p_{\tilde{\mathcal{R}}}(\alpha) \vee c_{\mathcal{R}}(\alpha) \neq c_{\tilde{\mathcal{R}}}(\alpha)) \wedge c_{\mathcal{R}}(\alpha) > \omega_{\tilde{\mathcal{R}}}(G, I). \end{cases}$$

- **Convergence.** Since the payment function is such that any move of an agent does not affect the cost of any other agent and has a codomain of cardinality 3, any agent can perform at most 3 moves during the game, thus the game converges in at most $3|I|$ moves.
- **Uniqueness.** We now show that the only Nash equilibrium for \mathcal{G} is $\tilde{\mathcal{R}}$. Suppose, by contradiction, that in the routing \mathcal{R} at Nash equilibrium, there exists an agent α such that $price_{\mathcal{R}}(\alpha) = 2$. Agent α can clearly perform a move by switching to a color $c_{\mathcal{R}}(\alpha) > \omega_{\tilde{\mathcal{R}}}(G, I)$, so as to pay a cost equal to 1, and this contradicts the hypothesis that \mathcal{R} is a Nash equilibrium. Analogously, suppose that in the routing \mathcal{R} at Nash equilibrium there exists at least one agent α such that $price_{\mathcal{R}}(\alpha) = 1$. In \mathcal{R} , α is thus assigned the color $c_{\mathcal{R}}(\alpha) > \omega_{\tilde{\mathcal{R}}}(G, I)$. Having already proven that α cannot pay a cost equal to 2, we have that no agent in the routing \mathcal{R} is assigned a color $c_{\mathcal{R}}(\alpha) \leq \omega_{\tilde{\mathcal{R}}}(G, I)$ unless $R(\alpha) = \tilde{R}(\alpha) \wedge c_{\mathcal{R}}(\alpha) = c_{\tilde{\mathcal{R}}}(\alpha)$, that is, unless he is using the same path and the same color of $\tilde{\mathcal{R}}$. Thus, agent α can perform a move by choosing the same path and color of $\tilde{\mathcal{R}}$, so as to pay 0, and this contradicts the hypothesis that \mathcal{R} is a Nash equilibrium. We have shown that in a routing \mathcal{R} at Nash equilibrium $price_{\mathcal{R}}(\alpha) = 0 \forall \alpha \in I$; by the definition of $price_{\mathcal{R}}$, each agent uses the same path and color of the routing $\tilde{\mathcal{R}}$, thus $\mathcal{R} = \tilde{\mathcal{R}}$.

□

Corollary 1. *Let \mathcal{A} be any k -approximation algorithm for the all-optical routing problem ($k = 1$ if \mathcal{A} is optimal). Then, there exists a payment function yielding a game \mathcal{G} , converging in at most $3|I|$ steps and having a price of anarchy equal to k . Moreover, the time complexity of the computation of the function is at most the same as \mathcal{A} .*

Proof. The claim is an immediate consequence of Theorem 1: by letting $\tilde{\mathcal{R}}$ be the routing returned by algorithm \mathcal{A} , the price of anarchy of \mathcal{G} is simply the approximation ratio of \mathcal{A} , and finally, each agent, in order to compute the payment function, has to find the same path and color computed by the algorithm \mathcal{A} . □

In the sequel, we will focus on the minimal and intermediate information levels. Before proceeding with the presentation of our results, let us first introduce two families of

payment functions that provide a useful characterization of a class of convergent games with the worst possible price of anarchy. Given a generic move μ performed by an agent α from \mathcal{R}_{old} to \mathcal{R}_{new} , let A be the set of agents sharing at least one edge with α in \mathcal{R}_{old} and no edges with α in \mathcal{R}_{new} and B be the set of agents sharing at least one edge with α in \mathcal{R}_{new} . We denote with Π the set of payment functions, satisfying the following conditions:

$$\forall \beta \in A, price_{\mathcal{R}_{new}}(\beta) \leq price_{\mathcal{R}_{old}}(\beta) \tag{1}$$

$$\forall \beta \in B, price_{\mathcal{R}_{new}}(\beta) > price_{\mathcal{R}_{old}}(\beta) \Rightarrow price_{\mathcal{R}_{new}}(\beta) \leq price_{\mathcal{R}_{new}}(\alpha) \tag{2}$$

Roughly speaking, condition (1) means that all the players sharing at least an edge with α before her move and no edge after such a move does not increase their costs. Condition (2) means that given a generic player β sharing at least an edge with α after μ and increasing their cost after such a move, the cost of β after μ is at most the cost of α after μ . Such conditions will be useful for proving the convergence to equilibrium.

Theorem 2. *All the payment functions belonging to Π are convergent.*

Proof. Let Ψ_{old} be the sequence of costs $price_{\mathcal{R}_{old}}(\beta) \forall \beta \in I$ ordered in a non-decreasing way and let Ψ_{new} be the sequence of costs $price_{\mathcal{R}_{new}}(\beta) \forall \beta \in I$ ordered in the same way, that is, the new sequence of costs yielded by the execution of a move μ performed by the agent α . By the definition of the class Π , μ can increase the cost of any other agent to reach, at most, the value $price_{\mathcal{R}_{new}}(\alpha)$. Since it holds $price_{\mathcal{R}_{new}}(\alpha) \leq price_{\mathcal{R}_{old}}(\alpha)$, we obtain that the sequence Ψ_{new} is lexicographically smaller than Ψ_{old} . Notice that such a lexicographic order induces a total order among the nodes of the Nash dynamics. The sequence $\Psi_0 = \langle 0, 0, 0, \dots \rangle$ is the bottom of the lexicographic ordering, hence the thesis. \square

As can be easily seen, the number of steps needed to converge to an equilibrium may not be polynomial in the dimensions of the instance.

We now define a subclass of Π of payment functions inducing games \mathcal{G} , having a price of anarchy $\rho = \frac{|I|}{\omega(G,I)}$, which is clearly the worst possible, since any feasible solution for the problem uses at most $|I|$ colors, while $\omega(G, I)$ is the optimal solution.

We denote as $\Xi = \Xi' \cup \Xi''$ the class of payment functions satisfying at least one of the following conditions:

- Subclass Ξ' : $\forall \alpha \in I, price(\alpha)$ depends only on the color $c(\alpha)$ and $c_{\mathcal{R}}(\alpha) \geq c_{\mathcal{R}'}(\alpha) \Rightarrow price_{\mathcal{R}}(\alpha) \geq price_{\mathcal{R}'}(\alpha)$ for any \mathcal{R} and \mathcal{R}' ; that is, the cost for each agent depends only on its own color in the routing and any other agent α never gets a benefit in performing a move $(\alpha, p_{old}, color_{old}, p_{new}, color_{new})$ where $color_{new} \geq color_{old}$.
- Subclass Ξ'' : $\forall \alpha \in I, price(\alpha) = \max_{e \in p(\alpha)} g(\sigma(e))$ where $g : 2^W \rightarrow \mathbf{R}$ is such that for any sets of colors A and $B, A \setminus \{d\} \subseteq B \Rightarrow g(A) \leq g(B \cup \{d'\})$ with $d \in A, d' \geq d$ and $d' \notin B$. In other words, the image $g(B \cup \{d'\})$ according to g of a set $B \cup \{d'\}$ containing all the elements of another set A except for at most one element (d) , which has to be replaced in the first set by a strictly greater element d' , cannot be smaller than the image $g(A)$ of the second set.

Theorem 3. $\Xi \subseteq \Pi$, that is any function $price \in \Xi$ is convergent. Moreover, for each payment function $price \in \Xi$, there exists a game $\mathcal{G} = (G, I, price)$ having a price of anarchy $\rho = \frac{|I|}{\omega(G,I)}$.

Proof. We prove that for each payment function $price, price \in \Xi' \Rightarrow price \in \Pi$ and $price \in \Xi'' \Rightarrow price \in \Pi$. Let $(\alpha, p_{old}(\alpha), color_{old}, p_{new}(\alpha), color_{new})$ be a generic move performed by the agent α from the routing \mathcal{R}_{old} to the routing \mathcal{R}_{new} . If $price \in \Xi'$, then both the conditions (1) and (2) are satisfied since $\forall \beta \neq \alpha, \beta \in I, price_{\mathcal{R}_{new}}(\beta) = price_{\mathcal{R}_{old}}(\beta)$. If $price \in \Xi''$, then condition (1) is satisfied since, setting $d = d', d \in A, d \notin B$ and $C = B \cup \{d\}$, it holds that $A \subseteq C \Rightarrow A \setminus \{d\} \subseteq B \Rightarrow g(A) \leq g(C)$, thus, the costs of the

agents sharing no edges with α in \mathcal{R}_{new} cannot increase. Condition (2) is satisfied by the payment function $price(\alpha) = \max_{e \in p^\alpha} g(\sigma(e))$, since we are assured that whenever the cost of the agents now sharing an edge with α in \mathcal{R}_{new} increase, it can never exceed the cost $price_{\mathcal{R}_{new}}(\alpha)$.

In order to prove the result for the price of anarchy, let G be a ring of $|I|$ nodes $(v_1, \dots, v_{|I|})$, and $I = \{\{v_i, v_{(i+1) \bmod |I|}\} | i \in \{1, \dots, |I|\}\}$ be the communication instance. The routing $\mathcal{R}^* = (P_{\mathcal{R}^*}, c_{\mathcal{R}^*})$, where $\forall i \in \{1, \dots, |I|\}$, $P_{\mathcal{R}^*}(\{v_i, v_{(i+1) \bmod |I|}\}) = \{\{v_i, v_{(i+1) \bmod |I|}\}\}$ and $c_{\mathcal{R}^*}(\{v_i, v_{(i+1) \bmod |I|}\}) = 1$, is such that $\omega_{\mathcal{R}^*}(G, I) = 1$, thus $w(G, I) = 1$. If we consider the routing $\mathcal{R} = (p_{\mathcal{R}}, c_{\mathcal{R}})$ where $\forall i \in \{1, \dots, |I|\}$, $p_{\mathcal{R}}(\{v_i, v_{(i+1) \bmod |I|}\}) = \{\{v_j \bmod |I|, v_{(j+1) \bmod |I|}\} | j \in \{i+1, \dots, i+|I|-1\}\}$ and $c_{\mathcal{R}}(\{v_i, v_{(i+1) \bmod |I|}\}) = i$ we have that \mathcal{R} is, a Nash equilibrium by the definition of Ξ . First we note that in \mathcal{R} , $\sigma_{\mathcal{R}}(\{v_i, v_{(i+1) \bmod |I|}\}) = \{1, \dots, i-1, i+1, \dots, |I|\}$, $\forall i \in \{1, \dots, |I|\}$, that is, all the edges $\{v_i, v_{(i+1) \bmod |I|}\}$ of the ring are occupied by all the colors except for the color i . The proof then proceeds for cases. Depending on the particular subclass of $price$, we have:

- Subclass Ξ' . No agent $\{v_i, v_{(i+1) \bmod |I|}\}, i \in \{1, \dots, |I|\}$ gets a benefit in performing a move, because he cannot use a color smaller than i ;
- Subclass Ξ'' . No agent $\alpha = \{v_i, v_{(i+1) \bmod |I|}\}, i \in \{1, \dots, |I|\}$ gets a benefit in performing a move, because if he moved to the edge $\{v_i, v_{(i+1) \bmod |I|}\}$ which is the only remaining one in the path of $|I|-1$ edges thus changing his color, his new path (in the routing \mathcal{R}') would contain at least an edge occupied by the set of colors $B \cup \{d'\} | B \supseteq \sigma_{\mathcal{R}}(e) \setminus c_{\mathcal{R}}(\alpha) \forall e \in p_{\mathcal{R}}(\alpha), d' \geq c_{\mathcal{R}}(\alpha)$, which would result in $price_{\mathcal{R}'}(\alpha) \geq price_{\mathcal{R}}(\alpha)$.

Since $\omega_{\mathcal{R}}(G, I) = |I|$, we obtain that the price of anarchy $\rho \geq |I|$. The claim for $w(G, I) = 1$ comes observing that in any feasible routing \mathcal{R} , it holds that $\omega_{\mathcal{R}}(G, I) \leq |I|$. The claim for $w(G, I) > 1$ can be obtained by replicating each request $w(G, I)$ times. By so doing, we can construct, similarly as above, a routing in Nash equilibrium using $|I|$ colors. \square

4. Minimal and Intermediate Payment Functions

Even though the results obtained in the case of complete information are fully satisfactory, they have been obtained under the assumption that each agent has full knowledge of the instance and the routing: this assumption is very strong and in real world networks might not be true either due to technological or privacy policies reasons. For this reason, in the sequel, we will focus on payment functions based on a more restricted level of information. Let us first define a set of payment functions that can be computed under a minimal or intermediate information level. Given a routing \mathcal{R} , we first propose suitable cost functions defined on the edges that will be used as building blocks for the definition of the mentioned payment functions:

- $col(e, \alpha) = f(c(\alpha))$: the amount charged to α on the edge e is the cost, according to f , of the color he uses.
- $max(e, \alpha) = \max_{k \in \sigma_{\mathcal{R}}(e)} f(k)$: the amount charged to α on the edge e is the cost of the highest color used along e (considering also the other agents).
- $sum(e, \alpha) = \sum_{k \in \sigma_{\mathcal{R}}(e)} f(k)$: the amount charged to α on the edge e is the sum of the costs of all the colors used along e .
- $avmax(e, \alpha) = \max_{k \in \sigma_{\mathcal{R}}(e)} \frac{f(k)}{|\sigma_{\mathcal{R}}(e)|}$: the amount charged to α on the edge e is the cost of the highest color used along e , averaged or shared among all the agents traversing e .
- $avsum(e, \alpha) = \sum_{k \in \sigma_{\mathcal{R}}(e)} \frac{f(k)}{|\sigma_{\mathcal{R}}(e)|}$: the amount charged to α on the edge e is the sum of the costs of all the colors used along e , averaged on all the agents traversing e .

Starting from any edge cost function $cost$, it is possible to define the following payment functions:

- $max-cost(\alpha) = \max_{e \in p(\alpha)} cost(e, \alpha)$: the price asked to α is the maximum cost, according to $cost$, of an edge used by α .

- $sum-cost(\alpha) = \sum_{e \in p(\alpha)} cost(e, \alpha)$: the price asked to α is the sum of the costs of the edges used by α .

The combination of the introduced edge cost functions with the above two strategies, that is, maximization or summation, gives rise to ten possible payment functions. In all cases, since the function f is non-decreasing, agents have an incentive to choose small colors so as to possibly minimize the overall number of used colors. Notice that all the ten introduced payment functions are computable under an intermediate information level. Moreover, $max-col$, $sum-col$ and $max-max$ require only the minimal level, as $max-col(\alpha) = \max_{e \in p(\alpha)} col(e, \alpha) = f(c(\alpha))$, $sum-col(\alpha) = \sum_{e \in p(\alpha)} col(e, \alpha) = |p(\alpha)| \cdot f(c(\alpha))$ and $max-max(\alpha) = \max_{e \in p(\alpha)} \max_{k \in \sigma(e)} f(k) = \max_{k \in \bar{\sigma}(p(\alpha))} f(k)$. The following lemma shows how the families of functions defined in the previous section nicely characterize the class of games induced by them.

Lemma 1. *The payment functions $max-col$ and $max-max$ belong to the class Ξ .*

Proof. It can be easily verified that, by definition, the payment function $max-col$ belongs to Ξ' , while the payment function $max-max$ belongs to Ξ'' . \square

Moreover, a similar result holds for $sum-col$.

Theorem 4. *The payment function $sum-col$ belongs to Π , that is, it is convergent, but there exists a game $\mathcal{G} = (G, I, sum-col)$ having a price of anarchy $\rho = \frac{|I|}{w(G, I)}$.*

Proof. The payment function $sum-col$ belongs to Π , since it clearly satisfies conditions (1) and (2). We now show that the game induced by the function $sum-col$ on the instance $I = \underbrace{\{\{a, b\}\}}_{n \text{ times}}$ in the graph depicted in Figure 1 does not admit a Nash equilibrium. Let k be $\lceil f(n) \rceil + 1$. The routing $\mathcal{R}^* = (P_{\mathcal{R}^*}, c_{\mathcal{R}^*})$, where $\forall \{a, b\} \in I, |P_{\mathcal{R}^*}(\{a, b\})| = k$, i.e., each request is routed on a different chain of k edges, and $c_{\mathcal{R}^*}(\{a, b\}) = 1$ is such that $\omega_{\mathcal{R}^*}(G, I) = 1$, thus $w(G, I) = 1$. If we consider the routing $\mathcal{R} = (p_{\mathcal{R}}, c_{\mathcal{R}})$ where $\forall \{a, b\} \in I, |P_{\mathcal{R}}(\{a, b\})| = 1$, then there exists a request using at least color n . Since \mathcal{R} is an equilibrium, the thesis follows. \square

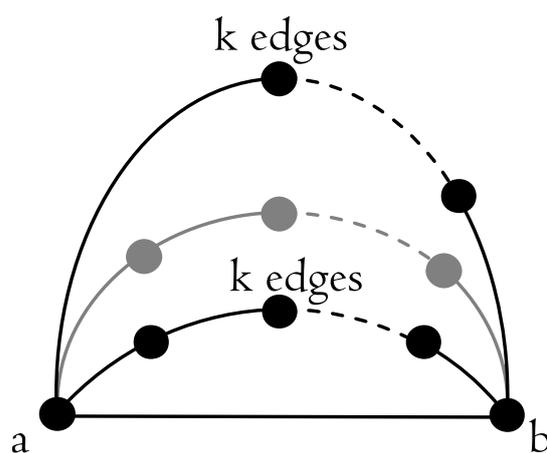


Figure 1. In this graph, nodes a and b are connected through a single edge and n chains of k edges.

Therefore, even if convergent, all the three minimal level functions yield the worst possible price of anarchy. Unfortunately, the following results show that, also, the remaining seven payment functions for the intermediate information level either are not convergent or yield the worst possible price of anarchy.

Theorem 5. *The payment function max-sum belongs to the class Ξ .*

Proof. We show that *max-sum* belongs to the subclass Ξ'' of Ξ . To this aim, it suffices to show that $g(\sigma(e)) = \text{sum}(e, \alpha) = \sum_{k \in \sigma(e)} f(k)$ satisfies, for any sets of colors A and B , the property $A \setminus \{d\} \subseteq B \Rightarrow g(A) \leq g(B \cup \{d'\})$ with $d \in A$, $d' \geq d$ and $d' \notin B$. Since $A \setminus \{d\} \subseteq B$ and f is positive and non decreasing, we have $g(B \cup \{d'\}) = f(d') + \sum_{k \in B} f(k) \geq f(d) + \sum_{k \in A \setminus \{d\}} f(k) = g(A)$. \square

Theorem 6. *Nash equilibria are not guaranteed to exist for the games induced by the payment functions.*

1. *sum-max when the pricing function f is unbounded;*
2. *max-avmax and sum-avmax when f is such that $\exists k : f(k) > 2f(1)$;*
3. *max-avsum and sum-avsum when the f is such that $\exists k : f(k) > f(1)$, that is, f is non-constant.*

The proof of Theorem 6 is deferred to the Appendix A. Finally, we have the following theorem concerning *sum-sum*.

Theorem 7. *The payment function sum-sum is not convergent when the pricing function f is such that $\exists k : f(k) > f(1) + f(2)$.*

Proof. We show that, for the problem defined by the graph depicted in Figure 2 and by the instance $I = \{\alpha = \{a, c\}, \beta = \{d, c\}\}$, the Nash dynamics are cyclic for any pricing function f , such that $\exists k : f(k) > f(1) + f(2)$.

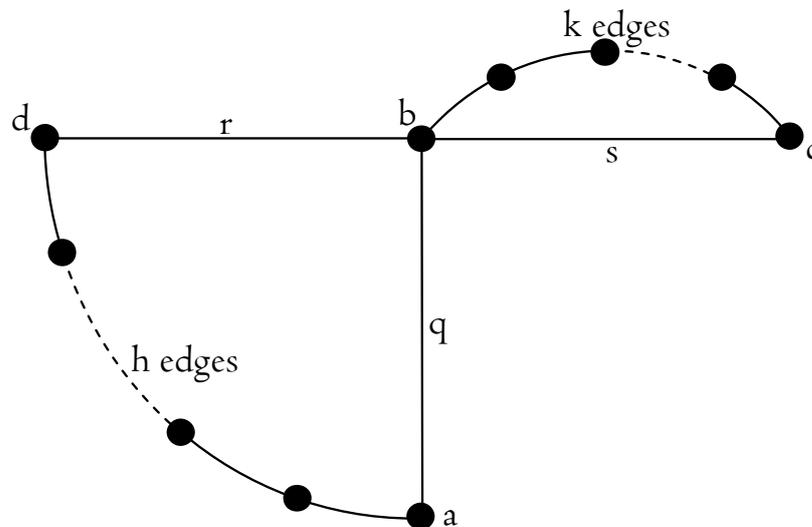


Figure 2. In this graph, there exists a path consisting of h edges between the nodes a and d , and there exist two paths, one consisting of a single edge and the other consisting of k edges between the nodes b and c .

The parameters n , h and k depend on the pricing function f and must be suitably tuned as shown in the sequel of the proof. Let q be the edge $\{a, b\}$, q' be the path obtained from the union of the path of h edges and the edge $\{d, b\}$, s be the edge $\{b, c\}$, s' be the path of k edges and r be the edge $\{d, b\}$.

The following conditions suffices to prove that the Nash dynamics are cyclic:

$$(k + 1)f(2) < 2f(n) + f(n + 1), \tag{3}$$

$$2f(n) + f(1) < (k + 1)f(2), \tag{4}$$

$$(h + 2)f(1) + f(2) < 2f(n + 1), \tag{5}$$

$$2f(n + 1) < (h + 2)f(1) + f(n). \tag{6}$$

Consider the initial routing $\mathcal{R}_0 = (R_0, c_0)$, displayed in Table 2.

Table 2. Routing \mathcal{R}_0 .

\mathcal{R}_0	R_0	c_0	$Sum\text{-}Sum_{\mathcal{R}_0}(\cdot)$
α	q, s	$n + 1$	$2f(n + 1) + f(n)$
β	r, s	n	$2f(n) + f(n + 1)$

At this first step, β plays the move $(\beta, \{r, s\}, n, \{r, s'\}, 2)$ because of condition (3), thus leading to the routing $\mathcal{R}_1 = (R_1, c_1)$, displayed in Table 3.

Table 3. Routing \mathcal{R}_1 .

\mathcal{R}_1	R_1	c_1	$Sum\text{-}Sum_{\mathcal{R}_1}(\cdot)$
α	q, s	$n + 1$	$2f(n + 1)$
β	r, s'	2	$(k + 1)f(2)$

Now, α plays the move $(\alpha, \{q, s\}, n + 1, \{q', s\}, 1)$ because of condition (5), thus leading to the routing $\mathcal{R}_2 = (R_2, c_2)$, displayed in Table 4.

Table 4. Routing \mathcal{R}_2 .

\mathcal{R}_2	R_2	c_2	$Sum\text{-}Sum_{\mathcal{R}_2}(\cdot)$
α	q', s	1	$(h + 2)f(1) + f(2)$
β	r, s'	2	$(k + 1)f(2) + f(1)$

β plays the move $(\beta, \{r, s'\}, 2, \{r, s\}, n)$ because of condition (4), thus leading to the routing $\mathcal{R}_3 = (R_3, c_3)$, displayed in Table 5.

Table 5. Routing \mathcal{R}_3 .

\mathcal{R}_3	R_3	c_3	$Sum\text{-}Sum_{\mathcal{R}_3}(\cdot)$
α	q', s	1	$(h + 2)f(1) + 2f(n)$
β	r, s	n	$2(f(n) + f(1))$

Finally, α plays the move $(\alpha, \{q', s\}, 1, \{q, s\}, n + 1)$ because of condition (6), thus leading the game back to the initial routing \mathcal{R}_0 (displayed in Table 2).

In order to complete the proof, we have to show how to satisfy the conditions (3)–(6) and how to tune the parameters n, k and h . The four conditions are satisfied by choosing k and h , such that

$$0 < \frac{2f(n) + f(1) - f(2)}{f(2)} < k < \frac{2f(n) + f(n + 1) - f(2)}{f(2)}$$

and

$$0 < \frac{2f(n + 1) - f(n) - 2f(1)}{f(1)} < h < \frac{2f(n + 1) - f(2) - 2f(1)}{f(1)}.$$

Choosing $n : f(n) > f(1) + f(2)$ assures that h and k always exist and are integers. \square

5. Results for Specific Topologies

In Section 4, we have shown that the price of anarchy for general topologies under

the intermediate and minimal levels of information can be very high. For this reason, always considering the intermediate and minimal levels of information, in this section, we consider networks having specific topologies, like chains (i.e., sites are connected along a line), and rings (i.e., cycles of sites) and trees. Let us first consider the minimal information level. Any strictly increasing payment function computable under this level, like $price(\alpha) = c(\alpha)$, induces games in which a routing \mathcal{R} at Nash equilibrium can be seen as a solution of the classical first-fit algorithm for the all-optical routing problem that assigns to each request the smallest available color. In particular, this solution is the one returned by first-fit when requests are considered in non-decreasing order of color in \mathcal{R} . Therefore, the induced price of anarchy is bounded by the approximation ratio of first-fit. Concerning the above-mentioned topologies, for any instance I , first-fit uses a number of colors that are at most $8\pi(G, I)$ in chains [24] and at most $\mathcal{O}((\log |I|)\pi(G, I))$ in trees [25]. Recalling that $\omega(G, I) \geq \pi(G, I)$, the following theorem holds.

Theorem 8. *The payment function $price(\alpha) = c(\alpha)$ induces games with a price of anarchy 8 in chains and $\rho = \mathcal{O}(\log |I|)$ in trees, both converging in $\omega_{\mathcal{R}}(G, I)^2$ steps from any initial routing.*

For rings, it is possible to use the result of [24] by routing requests on the chain obtained deleting an edge $e \in E$ of the ring. In fact, denoted as P , the path system containing all such connecting paths, for any routing \mathcal{R} with $P_{\mathcal{R}} = P$, it results in $\pi_{\mathcal{R}}(G, I) \leq 2\pi(G, I)$. In other words, this at most doubles the induced load. The following payment function forces Nash equilibria using the set of paths P :

$$price(\alpha) = \begin{cases} 1 & \text{if } e \in p(\alpha) \\ 1 - \frac{1}{c(\alpha)} & \text{otherwise} \end{cases}$$

Since the payment function belongs to Π and, when restricted to the routings \mathcal{R} with $P_{\mathcal{R}} = P$, is strictly increasing, it follows that each agent has an incentive in choosing the smallest available color. By recalling that, in this case, the load at most doubles with respect to the chain topology, the following theorem holds.

Theorem 9. *In ring networks, the above payment function induces games with price of anarchy 16 and converging in $\omega_{\mathcal{R}}(G, I)^2$ steps from any initial routing \mathcal{R} .*

In the remaining part of the section, we show how, raising the level of information to the intermediate one, it is possible to further reduce the price of anarchy for chains and rings. We first focus on ring topologies, as the corresponding results can be directly extended to chains. Our purpose is to force the agents to simulate the behavior of the online algorithm proposed by Slusarek [26]. In such an algorithm, the path system P is fixed and the optical spectrum is divided in shelves, numbered consecutively starting from 1. The color assigned to an arriving agent α is the lowest available one in the minimal shelf i , such that the load induced on the edges of α by all the agents up to shelf i is at most i . More precisely, denoted as $sh(w)$, the shelf of a given wavelength w , the load $l(\alpha, \mathcal{R})$ of an agent α according to a routing \mathcal{R} is

$$l(\alpha, \mathcal{R}) = \max_{e \in p(\alpha)} |\{w | w \in \sigma(e) \wedge sh(w) \leq sh(c(\alpha))\}|.$$

Clearly, any routing \mathcal{R} such that $P_{\mathcal{R}} = P$ has the same load $\pi_{\mathcal{R}}(G, I)$ and in the routing \mathcal{R} returned by the algorithm, at most $\pi_{\mathcal{R}}(G, I)$ shelves are used to allocate all the agents. Moreover, as shown in [26], during all the executions of the algorithm, the first shelf contains only one color and the other ones no more than three colors, thus yielding $\omega_{\mathcal{R}}(G, I) \leq 3 \cdot \pi_{\mathcal{R}}(G, I) - 2$ colors, that is, at most 3 times above the optimum.

In devising a payment function that mimics Slusarek’s algorithm, it is necessary to cope with several problems. First of all, it is necessary to fix a path system with a low

induced load. This is obtained by giving agents an incentive to use shortest paths, as in any routing \mathcal{R} satisfying such a property $\pi_{\mathcal{R}}(G, I) \leq 2\pi(G, I)$, that is, the induced load is at most doubled. Moreover, the convergence of the game is an issue, since the move of an agent might compromise the minimality of the payments of the agents in his shelf and in the above ones. Finally, in order for an agent to always find a proper shelf during the evolution of the game, we have to allow an unlimited number of colors to be contained in each shelf. Since, however, Nash equilibria will use only one wavelength of the first shelf and at most three in the other ones, we have to choose the function sh associating colors to shelves trying to minimize, for each i , the maximum third color of all the shelves up to i . To this aim, we partition the set of wavelengths W in two subsets W_1 and W_2 , in which $W_2 = \{2^i | i = 0, 1, 2, \dots\}$ is the set of the slack wavelengths. Each shelf has an infinite number of wavelengths and the first one of shelf 1 and the first three ones of all the other shelves correspond consecutively to small colors in W_1 , while the other ones to slack colors in W_2 , assigned to the different shelves according to a dove-tail technique in such a way that, for every $i \geq 1$ and $j > 3$, there exists a finite wavelength $w \in W_2$ that corresponds to the j -th color of shelf i . As can be easily checked, while the first color of shelf 1 is 3, the third color of each shelf $i > 1$ is at most $3i + \lfloor \log_2 3i \rfloor + 3$. Assuming that the function sh realizes the above mapping of colors and that n is the number of nodes in the ring, the payment function $price$ charges to agent α

$$\min \left\{ \left\lfloor \frac{n}{2} \right\rfloor, |p(\alpha)| + n \max\{0, l(\alpha, \mathcal{R}) - sh(c(\alpha))\} - \frac{1}{2 + sh(c(\alpha)) - \frac{1}{2c(\alpha)}} \right\}.$$

Intuitively, a routing \mathcal{R} is at Nash equilibrium if, and only if, each request uses a shortest path and the color that Slusarek’s algorithm assigns to him when agents arrive in non-decreasing order of color in \mathcal{R} , clearly not using colors in W_2 .

Theorem 10. *In ring networks, the game induced by the above payment function converges and has a price of anarchy $\rho = 6 + \mathcal{O}\left(\frac{\log(\pi(G,I))}{\pi(G,I)}\right)$.*

Proof. Let us consider any evolution of the game starting from any given configuration and let us show that it converges to a Nash equilibrium in a finite number of agent moves. At any intermediate configuration, let us partition I into two sets A and B , such that A contains all the agents that, in the sequel of the game, will never perform a move to a higher shelf. We show that, after a finite number of moves, a new agent will enter in A , so that in a finite number of steps, it will finally result $A = I$. This proves the convergence of the game, since starting from such final configuration, the agents cannot perform an infinite number of moves without raising their shelf.

First of all, we note that, since by construction, the number of wavelengths in each shelf is infinite, each agent α can always perform a move leading him in a feasible shelf i , that is, such that $i \geq l(\alpha, \mathcal{R})$. Moreover, the moves leading α on a shortest path and in a feasible shelf are the only ones which can make α ’s cost strictly smaller than $\lfloor \frac{n}{2} \rfloor$. Finally, an agent routing a shortest path and lying in a feasible shelf, maintaining his connection path, decreases his cost if he moves to a lower feasible shelf or to a smaller color of its shelf.

Consider a move of agent α_1 to a shelf i_1 . The proof proceeds by cases: after the move

1. α_1 will never perform a move increasing his shelf. In this case, α_1 is entered in A .
2. At a certain point of the game, α_1 performs a move increasing his shelf. This is due to another agent α_2 who has performed a move to a shelf $i_2 \leq i_1$. However, it cannot be $i_2 = i_1$, as α_2 increases α_1 ’s load to $i_1 + 1$ and would have the same load $i_1 + 1$ in shelf i_1 (i.e., i_1 is not feasible for α_2), thus not decreasing his cost. We can then continue applying the same analysis to α_2 . Since the number of shelves in which agents can

move is bounded by $\pi(G, I)$, we must finally arrive to an agent α_j with $j \leq \pi(G, I)$ for which the first case holds.

As already observed, any routing \mathcal{R} at Nash equilibrium corresponds to the output of Slusarek's algorithm without using colors in W_2 when the agents arrive in non decreasing order of color in \mathcal{R} . Thus, the maximum used color in \mathcal{R} is at most the third one of shelf $\pi(G, I)$ if $\pi(G, I) > 1$, while it is equal to the first one of shelf 1 otherwise. This shows that the maximum used color is at most $3\pi_{\mathcal{R}}(G, I) + \lfloor \log_2 3\pi_{\mathcal{R}}(G, I) \rfloor + 3$. Since \mathcal{R} uses shortest paths, $\pi_{\mathcal{R}}(G, I) \leq 2\pi(G, I)$ and the resulting price of anarchy is $\rho = 6 + \mathcal{O}\left(\frac{\log(\pi_{\mathcal{R}}(G, I))}{\pi_{\mathcal{R}}(G, I)}\right)$. \square

Clearly, a similar function

$$price' = \min \left\{ n, 1 + 2n \max\{0, load(\alpha, \mathcal{R}) - sh(c(\alpha))\} - \frac{1}{2 + sh(c(\alpha)) - \frac{1}{2c(\alpha)}} \right\}$$

can be used also in chains, for which connection paths are unique and thus no doubling of the load occurs.

Theorem 11. *In chain networks, the game induced by the above payment functions converges and has a price of anarchy $\rho = 3 + \mathcal{O}\left(\frac{\log(\pi(G, I))}{\pi(G, I)}\right)$.*

6. Conclusions

We have considered the problem of determining suitable payment functions for the non-cooperative agents of an all-optical network, in order to induce Nash equilibria using a low number of wavelengths. In particular, we have identified three levels of information, specifying the local knowledge that the agent can exploit when computing their payments. While the complete information level has been fully understood, under the lower levels, the main left open question is the determination of functions that on every topology yield Nash equilibria with a performance better than the worst possible one assigning a different color to each agent. Moreover, still under incomplete information, it would be also nice to improve and extend our results on specific topologies. Furthermore, given that a lot of literature, as discussed in the introduction, focuses on the minimization of the cost due to hardware equipment in optical networks, it would be nice to study the characteristics and performance of games combining our measure (the number of used wavelengths) with other kinds of hardware costs. Moreover, since our findings outline nice connections between payment functions and online algorithms, that allow one to cope with the arbitrary order of the moves of the agents, it would be nice to understand the conditions and eventual systematic methods allowing one to get converging payment functions preserving online algorithms performances under incomplete information. Finally, we would like to remark that this paper constitutes a theoretical contribution to the study of optical networks from a decentralized multi-agent point of view, and it would be interesting to validate the obtained results and proposed approach by some simulation studies to be performed in a future work. To this respect, we believe that many worse lower bounds to the price of anarchy are achieved by exploiting specific and pathological instances, and therefore, the price of anarchy may turn out to be much lower, in practice, in the cases for which the existence of instances with a very large price of anarchy has been theoretically proven.

Author Contributions: Conceptualization, V.B., M.F. and L.M.; methodology, V.B., M.F. and L.M.; formal analysis, V.B., M.F. and L.M.; writing—original draft preparation, V.B., M.F. and L.M.; writing—review and editing, V.B. and L.M.; funding acquisition, V.B., M.F. and L.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially supported by the PRIN 2008 research project COGENT (Computational and Game-Theoretic Aspects of Uncoordinated Networks), funded by the Italian Ministry of University and Research.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A. Proof of Theorem 6

Proof. Claim 1. We show that the game induced by the function *sum-max* on the instance $I = \{\{a, c_i\}, \underbrace{\{b, d_i\}}_{m \text{ times}} \mid i \in \{1, 2, \dots, n, n + 1\}\}$ in the graph depicted in Figure A1, does not admit a Nash equilibrium if f is unbounded. In such a graph, there exist two paths between the nodes a and b , one consisting of a unique edge and the other one having h edges. Similarly, there exist two paths between any pair of nodes b and c_i having one and k edges, respectively. The parameters m, n, h and k depend on f and must be carefully tuned, as shown in the sequel.

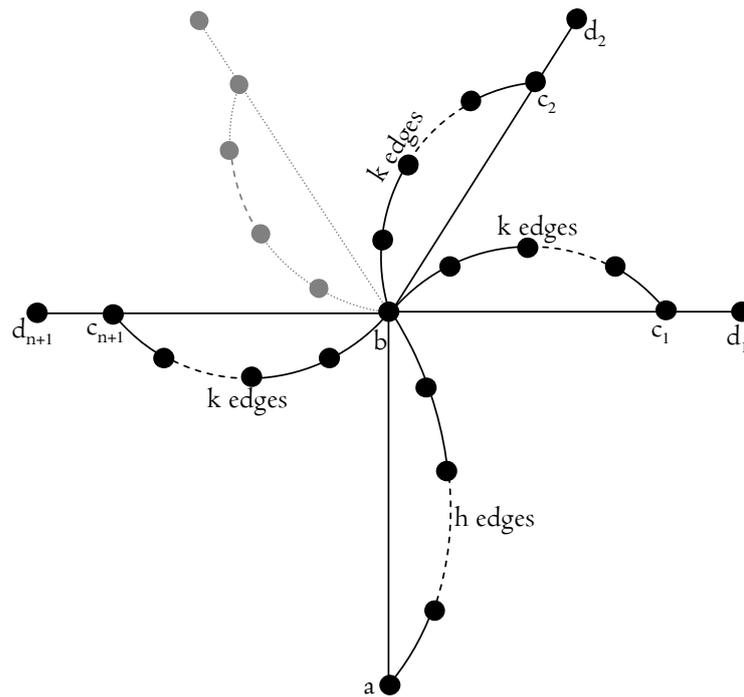


Figure A1. The graph not yielding a Nash equilibrium for the payment function *sum-max*.

We now give a set of conditions which is sufficient to prove the non-existence of a Nash equilibrium. First of all, we have three technical conditions:

$$n > m > 1, \tag{A1}$$

$$f(n + 1) > f(n), \tag{A2}$$

$$f(m + 1) > f(m). \tag{A3}$$

Clearly, in any Nash equilibrium, the agents $\{b, d_i\}$ can use at most the color $m + 1$ (or a color x such that $f(x) = f(m + 1)$), if the edge $\{b, c_i\}$ is used by an agent $\{a, c_i\}$ with a color at most m , and can use at most color m otherwise. Because of technical condition (A2) and by imposing the following conditions,

$$2f(n) < (h + k)f(1), \tag{A4}$$

$$2f(n) < hf(1) + f(m + 1), \tag{A5}$$

$$(k + 1)f(m) < 2f(n + 1), \tag{A6}$$

$$(k + 1)f(m) < (h + k)f(1), \tag{A7}$$

$$(k + 1)f(m) < hf(1) + f(m + 1), \tag{A8}$$

in a Nash equilibrium exactly n agents $\{a, c_i\}$ must use the edge $\{a, b\}$ with colors ranging from 1 to n . In the right side of (A5) and (A8), $f(m + 1)$ replaces $f(1)$, since we impose the following condition, by which the m agents $\{b, d_i\}$ have to use the edge $\{b, c_i\}$ if they can use a color at most $m + 1$:

$$2f(m + 1) < (k + 1)f(1). \tag{A9}$$

In order to explain how exactly n agents $\{a, c_i\}$ must use the edge $\{a, b\}$ with colors ranging from 1 to n , we proceed by cases.

- If a color $x \leq m$ is free on the edges $\{a, b\}$ and $\{b, c_i\}$, the request $\{a, c_i\}$ is not in equilibrium using a color $y \geq n + 1$ or using the h edges chain (see conditions (A4), (A5) and (A9)).
- If a color $x \leq m$ is free on the edge $\{a, b\}$ and is not on the one $\{b, c_i\}$ (i.e., x is free on the corresponding k edges chain), the request $\{a, c_i\}$ is not in equilibrium using a color $y \geq n + 1$ or using the h edges chain (see conditions (A6)–(A9)).
- If a color x such that $m + 1 \leq x \leq n$ is free on the edge $\{a, b\}$, the request $\{a, c_i\}$ is not in equilibrium using a color $y \geq n + 1$ or using the h edges chain, since in a Nash equilibrium, the agents $\{b, d_i\}$ cannot use the edge $\{b, c_i\}$ with color x (see conditions (A4), (A5) and (A9)).

Now let $\alpha = \{a, c_i\}$ be the agent not involved in the previous analysis, and β_1, \dots, β_m the m agents $\{b, c_i\}$. α is not in equilibrium using the k edges chain by imposing

$$2f(n + 1) < (h + k)f(1), \tag{A10}$$

$$2f(n + 1) < (1 + k)f(n + 1). \tag{A11}$$

α is not in equilibrium using the k edges chain with a color $y \geq m + 1$ by imposing

$$2f(n + 1) < (h + 1)f(m + 1). \tag{A12}$$

If the edge $\{b, c_i\}$ is free, α is not in equilibrium using the edges $\{a, b\}$ and $\{b, c_i\}$ where it can use a color $y \geq n + 1$ by imposing

$$(h + 1)f(1) < 2f(n + 1). \tag{A13}$$

If the edge $\{b, c_i\}$ is used by the β agents, α is not in equilibrium using the h edges chain and the edge $\{b, c_i\}$, neither with color 1, by imposing

$$2f(n + 1) < hf(1) + f(m + 1). \tag{A14}$$

Moreover, because of condition (A9), the β agents are not in equilibrium using the k edges chain if all the colors $y \geq m + 1$ are free on the edge $\{b, c_i\}$, and are not in equilibrium using the unique edge $\{b, c_i\}$ if on it a color greater or equal to $n + 1$ is present by imposing

$$kf(m) < f(n + 1). \tag{A15}$$

We now prove that the agents α and β_j can never be at equilibrium proceeding by cases.

- The agent α uses the chain of h edges with a color at most equal to m (by condition (A12)), and there exists at least one agent β_j using the chain of k edges. This is not an equilibrium, since by condition (A9) β_j can move to the edge $\{b, c_i\}$.
- The agent α uses the chain of h edges with a color at most equal to m (by condition (A12)), and all the agents β use the edge $\{b, c_i\}$. This is not an equilibrium, since by condition (A14), α can move on the path formed by two edges using the color $n + 1$.
- The agent α uses the edge $\{a, b\}$ with a color at least equal to $n + 1$ and there exists at least one agent β_j using the unique edge $\{b, c_i\}$.

This is not an equilibrium, since by condition (A15), β_j can move to the chain of k edges.

- The agent α uses the edge $\{a, b\}$ with a color at least equal to $n + 1$ and all the agents β use the chain of k edges.

This is not an equilibrium since by condition (A13), α can move to the chain of h edges.

We conclude the proof by showing how the imposed conditions can be always satisfied.

We first note that some conditions are implied by other ones:

$$(A14) \wedge (A1) \wedge (A3) \Rightarrow (A12)$$

$$\text{since } 2f(n+1) < hf(1) + f(m+1) < hf(m+1) + f(m+1) = (h+1)f(m+1);$$

$$(A14) \Rightarrow (A5) \text{ since } 2f(n) \leq 2f(n+1) < hf(1) + f(m+1);$$

$$(A9) \wedge (A14) \Rightarrow (A10) \text{ since } 2f(n+1) < hf(1) + f(m+1) < hf(1) + kf(1) = (h+k)f(1);$$

$$(A10) \Rightarrow (A4) \text{ since } 2f(n) \leq 2f(n+1) < (h+k)f(1);$$

$$(A9) \wedge (A1) \wedge (A3) \Rightarrow (A11) \text{ since } k > \frac{f(m+1)}{f(1)} > 1 \Rightarrow 2f(n+1) < (1+k)f(n+1);$$

$$(A15) \Rightarrow (A6) \text{ since } (k+1)f(m) = kf(m) + f(m) < f(n+1) + f(n+1) = 2f(n+1);$$

$$(A6) \wedge (A10) \Rightarrow (A7) \text{ since } (k+1)f(m) < 2f(n+1) < (h+k)f(1);$$

$$(A6) \wedge (A14) \Rightarrow (A7) \text{ since } (k+1)f(m) < 2f(n+1) < hf(1) + f(m+1).$$

From the remaining conditions (A9), (A13), (A14) and (A15) we obtain

$$\frac{2f(n+1) - f(m+1)}{f(1)} < h < \frac{2f(n+1) - f(1)}{f(1)}$$

and

$$\frac{2f(m+1) - f(1)}{f(1)} < k < \frac{f(n+1)}{f(m)}.$$

The existence of h and k integers is guaranteed by $f(m+1) > 2f(1)$ and $f(n+1) > \frac{2f(m)f(m+1)}{f(1)}$, which make the width of the above intervals greater than 1.

Finally, in order to meet all the conditions including the initial technical ones, it is sufficient to choose

$$m \text{ such that } \begin{cases} f(m+1) > 2f(1), \\ f(m+1) > f(m), \\ m > 1. \end{cases}$$

and

$$n \text{ such that } \begin{cases} f(n+1) > \frac{2f(m)f(m+1)}{f(1)}, \\ f(n+1) > f(n), \\ n > m. \end{cases}$$

Since f is unbounded, all the constraints are satisfiable.

Claim 2. We show that, for the problem defined by the graph depicted in Figure A2 and by the instance

$$I = \{\{a, c_i\}, \{b, c_i\} \mid i \in \{1, 2, \dots, n\}\}$$

no Nash equilibrium exists for any pricing function f , such that $\exists k : f(k) > 2f(1)$.

The parameter n depends on f and must be suitably tuned. We set n as the smallest integer for which $f(n) > 2f(1)$. In any routing, there must exist an agent $\{a, c_i\}$ using a color at least equal to n . Let $\alpha = \{a, c_i\}$ be one of these agents and $\beta = \{b, c_i\}$. Clearly, in any routing at Nash equilibrium, β must use a color strictly less than n . We prove the claim by considering all the possible routings.

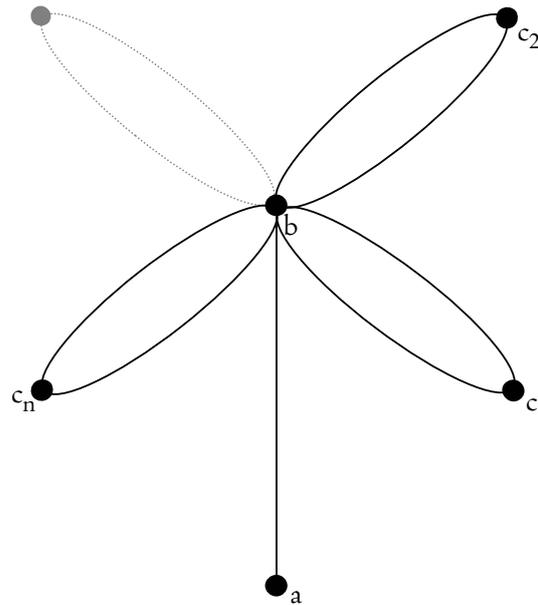


Figure A2. In this graph, there exist two edges between any pair of nodes b and c_i .

- α and β choose the same edge.
This is not an equilibrium, since β can always play a move so as not to share an edge with α . This is assured by the condition

$$f(1) < \frac{f(n)}{2}.$$

We note that from β 's point of view, the two pricing functions *max-avsum* and *sum-avsum* are essentially the same, since all the possible paths for β consist of a single edge.

- α and β choose different edges.
This is not an equilibrium, since α can always play a move so as to share an edge with β . This is assured by the condition

$$\frac{f(n)}{2} < f(n).$$

We note that, in a routing at Nash equilibrium, the edge $\{a, b\}$ has a cost of $\frac{f(n)}{n}$ for each agent $\{a, c_i\}$. For the pricing function *max-avmax*, the above move is trivially an improving one for α , while for the pricing function *sum-avmax*, this is assured by the condition $\frac{f(n)}{n} < f(n)$.

Claim 3. We show that for the problem defined by the graph depicted in Figure A2 and by the instance

$$I = \{\{a, c_i\}, \{b, c_i\} | i \in \{1, 2, \dots, n\}\}$$

no Nash equilibrium exists for any pricing function f such that $\exists k : f(k) > f(1)$. The parameter n depends on f and must be suitably tuned. We set n as the smallest integer for which $f(n) > f(1)$. In any routing, there must exist an agent $\{a, c_i\}$ using a color at least equal to n . Let $\alpha = \{a, c_i\}$ be such an agent and let $\beta = \{b, c_i\}$. Clearly, in any routing at Nash equilibrium, β must use a color strictly less than n . We prove the claim by considering all the possible routings.

- α and β choose the same edge.

This is not an equilibrium, since β can always play a move so as not to share an edge with α . This is assured by the condition

$$f(1) < \frac{f(1) + f(n)}{2}.$$

We note that, from β 's point of view, the two pricing functions *max-avsum* and *sum-avsum* are essentially the same, since all the possible paths for β consist of a single edge.

- α and β choose different edges.

This is not an equilibrium, since α can always play a move so as to share an edge with β . This is assured by the condition

$$\frac{f(1) + f(n)}{2} < f(n).$$

We note that, in a routing at Nash equilibrium, the edge $\{a, b\}$ has a cost of $\frac{(n-1)f(1)+f(n)}{n}$ for each agent $\{a, c_i\}$. For the pricing function *max-avsum*, the above move is trivially an improving one for α , while for the pricing function *sum-avsum*, this is assured by the condition $\frac{(n-1)f(1)+f(n)}{n} < f(n)$.

□

References

1. Brackett, C.A. Dense wavelength division multiplexing networks: Principles and applications. *IEEE J. Sel. Areas Commun.* **1990**, *8*, 948–964. [[CrossRef](#)]
2. Beauquier, B.; Bermond, J.C.; Gargano, L.; Hell, P.; Perennes, S.; Vaccaro, U. Graph Problems Arising from Wavelength-Routing in All-Optical Networks. In Proceedings of the 2nd Workshop on Optics and Computer Science (Part of IPPS'97), Geneva, Switzerland, 1 April 1997.
3. Gargano, L.; Vaccaro, U. Routing in All-Optical Networks: Algorithmic and Graph-Theoretic Problems. In *Numbers, Information and Complexity*; Kluwer Academic: Amsterdam, The Netherlands, 2000; pp. 555–578.
4. Fabrikant, A.; Luthra, A.; Maneva, E.; Papadimitriou, C.H.; Shenker, S. On a network creation game. In Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC), Boston, MA, USA, 13–16 July 2003; pp. 347–351.
5. Fabrikant, A.; Papadimitriou, C.H.; Talwar, K. The complexity of pure equilibria. In Proceedings of the 36th ACM Symposium on Theory of Computing (STOC), Chicago, IL, USA, 13–15 June 2004; ACM Press: New York, NY, USA, 2004; pp. 604–612.
6. Fotakis, D.; Kontogiannis, S.C.; Koutsoupias, E.; Mavronicolas, M.; Spirakis, P.G. The structure and complexity of Nash equilibria for a selfish routing game. *Theor. Comput. Sci.* **2009**, *410*, 3305–3326. [[CrossRef](#)]
7. Milchtaich, I. Congestion games with player-specific payoff functions. *Games Econ. Behav.* **1996**, *13*, 111–124. [[CrossRef](#)]
8. Owen, G. *Game Theory*, 3rd ed.; Academic Press: Cambridge, MA, USA, 1995.
9. Papadimitriou, C.H.; Yannakakis, M. On complexity as bounded rationality. In Proceedings of the 26th Annual ACM Symposium on the Theory of Computing (STOC), Montreal, QC, Canada, 23–25 May 1994; pp. 726–733.
10. Rosenthal, R.W. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* **1973**, *2*, 65–67. [[CrossRef](#)]
11. Vetta, A. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS), Vancouver, BC, Canada, 16–19 November 2002; pp. 416–425.
12. Koutsoupias, E.; Papadimitriou, C. Worst-case equilibria. In Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS), Trier, Germany, 4–6 March 1999; Volume 1563, pp. 387–396.
13. Mavronicolas, M.; Spirakis, P. The price of selfish routing. In Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing (STOC), Crete, Greece, 6–8 July 2001; pp. 510–519.
14. Roughgarden, T.; Tardos, E. How bad is selfish routing? *J. ACM* **2002**, *49*, 236–259. [[CrossRef](#)]
15. Nisan, N.; Roughgarden, T.; Tardos, E.; Vazirani, V.V. *Algorithmic Game Theory*; Cambridge University Press: Cambridge, UK, 2007.
16. Georgakopoulos, G.F.; Kavvadias, D.J.; Sioutis, L.G. Nash equilibria in all-optical networks. *Discret. Math.* **2009**, *309*, 4332–4342. [[CrossRef](#)]
17. Mertikopoulos, P.; Moustakas, A.L.; Tzanakaki, A. Boltzmann meets Nash: Energy-efficient routing in optical networks under uncertainty. *arXiv* **2016**, arXiv:1605.01451.
18. Pavel, L. *Game Theory for Control of Optical Networks*; Birkhäuser: Basel, Switzerland, 2012.
19. Flammini, M.; Marchetti-Spaccamela, A.; Monaco, G.; Moscardelli, L.; Zaks, S. On the complexity of the regenerator placement problem in optical networks. *IEEE/ACM Trans. Netw.* **2011**, *19*, 498–511. [[CrossRef](#)]

20. Flammini, M.; Monaco, G.; Moscardelli, L.; Shachnai, H.; Shalom, M.; Tamir, T.; Zaks, S. Minimizing total busy time in parallel scheduling with application to optical networks. *Theor. Comput. Sci.* **2010**, *411*, 3553–3562. [[CrossRef](#)]
21. Flammini, M.; Monaco, G.; Moscardelli, L.; Shalom, M.; Zaks, S. Optimizing regenerator cost in traffic grooming. *Theor. Comput. Sci.* **2011**, *412*, 7109–7121. [[CrossRef](#)]
22. Flammini, M.; Monaco, G.; Moscardelli, L.; Shalom, M.; Zaks, S. On the Complexity of the Regenerator Cost Problem in General Networks with Traffic Grooming. *Algorithmica* **2014**, *68*, 671–691. [[CrossRef](#)]
23. Flammini, M.; Monaco, G.; Moscardelli, L.; Shalom, M.; Zaks, S. Selfishness, collusion and power of local search for the ADMs minimization problem. *Comput. Netw.* **2008**, *52*, 1721–1731. [[CrossRef](#)]
24. Pemmaraju, S.V.; Raman, R.; Varadarajan, K. Max-Coloring and Online Coloring with Bandwidths on Interval Graphs. *ACM Trans. Algorithms* **2011**, *7*, 35:1–35:21. [[CrossRef](#)]
25. Bartal, Y.; Leonardi, S. On-line Routing in All-Optical Networks. *Theor. Comput. Sci. Spec. Issue* **1999**, *221*, 19–39. [[CrossRef](#)]
26. Slusarek, M. Optimal On-Line Coloring of Circular Arc Graphs. *Inform. Theorique Appl.* **1995**, *29*, 423–429. [[CrossRef](#)]