Selfishness, Collusion and Power of Local Search for the ADMs Minimization Problem ⋆

Michele Flammini [a] Gianpiero Monaco [a,∗] Luca Moscardelli [a]
Mordechai Shalom [b] Shmuel Zaks [c]

[a] *Università degli Studi dell'Aquila, L'Aquila, Italy*

[b] *TelHai Academic College, Upper Galilee, 12210, Israel*

[c] *Department of Computer Science, Technion, Haifa, Israel*

Abstract

We consider non cooperative games in all-optical networks where users share the cost of the used ADM switches for realizing given communication patterns. We show that the two fundamental cost sharing methods, Shapley and Egalitarian, induce polynomial converging games with price of anarchy at most $\frac{5}{3}$, regardless of the network topology. Such a bound is tight even for rings. Then, we show that if collusion of at most k players is allowed, the Egalitarian method yields polynomially converging games with price of collusion between $\frac{3}{2}$ and $\frac{3}{2}+\frac{1}{k}$. This result is very interesting and quite surprising, as the best known approximation ratio, that is $\frac{3}{2}+\epsilon$, can be achieved in polynomial time by uncoordinated evolutions of collusion games with coalitions of increasing size. Finally, the Shapley method does not induce well defined collusion games, but can be exploited in the definition of local search algorithms with local optima arbitrarily close to optimal solutions. This would potentially generate PTAS, but unfortunately the arising algorithm might not converge. The determination of new cost sharing methods or local search algorithms reaching a compromise between Shapley and Egalitarian is thus outlined as being a promising and worth pursuing investigating direction.

Key words: Optical Networks, Add-Drop Multiplexer (ADM), Nash Equilibria, Price of Anarchy, Price of Collusion.

⋆ This research was partly supported by the EU Project "Graphs and Algorithms in Communication Networks (GRAAL)" - COST Action TIST 293.

∗ Corresponding author. Tel: +39 0862433735; Fax: +39 0862433131; Address: Department of Computer Science Via Vetoio Coppito 67100 L'Aquila, Italy.

Email addresses: flammini@di.univaq.it (Michele Flammini), gianpiero.monaco@di.univaq.it (Gianpiero Monaco), moscardelli@di.univaq.it (Luca Moscardelli), cmshalom@telhai.ac.il

Preprint submitted to Computer Networks

1 Introduction

1.1 Background

All-optical networks have been largely investigated in recent years due to the promise of data transmission rates several orders of magnitudes higher than current networks [6,7,22,24]. Major applications are in video conferencing, scientific visualization and real-time medical imaging, high-speed supercomputing and distributed computing [11,22].

The key to high speeds in all-optical networks is to maintain the signal in optical form, thereby avoiding the prohibitive overhead of conversion to and from the electrical form at the intermediate nodes. The high bandwidth of the optical fiber is utilized through *wavelength-division multiplexing*: two signals connecting different source-destination pairs may share a link, provided they are transmitted on carriers having different wavelengths (or colors) of light. The optical spectrum being a scarce resource, given communication patterns in different topologies are often designed so as to minimize the total number of used colors, also as a comparison with the trivial lower bound provided by maximum load, that is the maximum number of connection paths sharing a same physical edge [3,19].

When the various parameters comprising the switching mechanism in these networks became clearer, the focus of studies shifted, and today a large portion of research concentrates with the total hardware cost. This is modelled by considering the basic electronic switching units of the electronic Add-Drop Multiplexer (ADM) and focusing on the total number of these hardware components. Each lightpath uses two ADMs, one at each endpoint. If two non-overlapping lightpaths are assigned the same wavelength and are incident to the same node, then they can use the same ADM. Thus, an ADM may be shared by at most two lightpaths. The problem of minimizing the number of ADMs was introduced in [21] for ring networks. For such a topology it was shown to be NP-complete [13] and an approximation algorithm with approximation ratio $3/2$ was presented in [9] and improved in [31,14] to $10/7 + \epsilon$ and $10/7$ respectively. For general topologies an algorithm with approximation ratio $8/5$ is presented in [13]. For the same problem, algorithms with approximation ratio $3/2 + \epsilon$ were provided in [8,17].

In a distributed and decentralized environment characterizing an optical communication network, besides the classical design of centralized algorithms optimizing the resources utilization, the analysis of the uncooperative interaction between the network users and the design of distributed algorithms call for

(Mordechai Shalom), zaks@cs.technion.ac.il (Shmuel Zaks).

more research effort. On this respect, Game Theory and the associated concept of Nash equilibria [28] have recently emerged as a powerful tool for modelling and analyzing such a lack of coordination. In this setting, each communication request is handled by an agent (or player) selfishly performing *moves*, i.e. changing her routing strategy in order to maximize her own benefit. A Nash equilibrium is a solution of the game in which no agent gains by unilaterally changing her routing strategy. If Nash equilibria are reached in a polynomial number of selfish moves, and finding an improving user move is a problem solvable in polynomial time, such an uncooperative process naturally defines a distributed polynomial time algorithm. However, due to the lack of cooperation among the players, Nash equilibria are known not to always optimize the overall performance. Such a loss [10,1] has been formalized by the so-called *price of anarchy* (resp. *optimistic price of anarchy*), defined as the ratio between the cost of the worst (resp. best) Nash equilibrium and the one of an optimal centralized solution.

There exists a vast literature on Nash Equilibria in communication networks [25,29]. The problem of investigating the existence and performance of Nash equilibria in all-optical networks has been first considered [4,5] with respect to the minimization of the number of used wavelengths. In such a setting, a service provider has to satisfy a given set of point-to-point communication requests, charging each of them a cost depending on its wavelength and on the wavelengths of the other requests met along its path in the network. Each request is issued by a non-cooperative agent interested only in the minimization of her own cost. In a similar setting [20], the authors focus on the complexity of recognizing and computing Nash equilibria.

Even if in non-cooperative games players are usually considered to act selfishly and independently, an interesting investigated issue is the one of collusion. Roughly speaking, collusion allows two or more players forming a coalition to come to an agreement in order to obtain a gain by changing at the same time their strategies. In this framework, a Nash equilibrium is a solution in which there exists no coalition of players having convenience in changing their strategies. The lack of performance with respect to the optimal solution has been measured by the *price of collusion* [18,23], where the authors focused on a particular class of games, the congestion games, assuming the players partitioned into sets of coalitions. Earlier [2,26], the authors provided other related equilibrium concepts, such as strong or coalition-proof equilibria, ensuring that coalitions have no incentive to form. Coalitions have also been considered from the perspective of the algorithmic mechanism design, with emphases on group-strategyproof mechanisms [27].

1.2 Our Contribution

Following the research direction outlined in [16], in this paper we are interested in analyzing the non-cooperative scenario in which the users of an optical network interact sharing the cost of the used hardware components. More precisely, we focus on ADM switches, considering the game in which their total cost is divided between the users according to two fundamental cost sharing methods: the Shapley [32] method, in which the agents using an ADM pay for it by equally splitting its cost, and the Egalitarian one, where the whole hardware cost is equally split among all the players.

We show that without collusion the two cost sharing methods are equivalent and induce games always convergent in polynomial time, i.e. the players always reach an equilibrium configuration within a polynomial number of selfish moves. Moreover, we prove that the arising price of anarchy is at most $\frac{5}{3}$ regardless of the network topology, and that such a result is tight even for rings. This result is very interesting, as it matches the performance of three different algorithms [13,8].

Under the assumption that the collusion of at most k players is allowed, only the Egalitarian cost sharing method yields a well-founded definition of induced game. We show that such a game is still convergent, and its price of collusion is $\frac{3}{2} + \frac{1}{k}$. This result is quite surprising, as the best known approximation ratio reached by a centralized algorithm [8], that is $\frac{3}{2} + \epsilon$, can be achieved in polynomial time by uncoordinated evolutions of collusion games with coalitions of increasing size. As already remarked, this has the additional appreciable effect of yielding a polynomial time approximation algorithm of distributed nature with the best so far achievable optimization performance.

Finally, always under the assumption of collusion, the Shapley method does not induce well defined games, but it can be exploited in the definition of proper neighborhoods in local search algorithms. The arising local optima are arbitrarily close to optimal solutions, that is at most $1 + \frac{2}{k}$ times the optimum, thus potentially generating distributed PTAS; unfortunately, the arising algorithms might not converge and such local optima might even not exist at all. However, this sheds some light on the effectiveness of local search in improving the current approximation factors. In fact, the determination of new cost sharing methods reaching a compromise between the Shapley and Egalitarian ones in terms of optimization and performance is thus outlined as a promising and worth pursuing investigating direction that will possibly capture future research attention.

The paper is organized as follows. In the next section we give the basic notation and definitions and show some preliminary results. In Section 3 and 4 we focus on the ADM minimization, and we show the results concerning Nash equilibria without and with collusion, respectively. Finally, in Section 5 we discuss the power of local search algorithm, give some conclusive remarks and discuss some open questions.

2 Model and Preliminary Results

An instance of the ADMs minimization problem is a pair (G, P), where G is an undirected graph and $P = \{p_1, \ldots, p_n\}$ is a multi-set of n simple paths in G, also called lightpaths or requests.

A coloring (or wavelength assignment) of (G, P) is a function $w : P \mapsto \mathbb{N}^+ = \{1, 2, \ldots\}$ such that $w(p_i) \neq w(p_j)$ for any pair of paths $p_i, p_j \in P$ sharing an edge in G.

Given a coloring function w, a valid cycle (resp. chain) is a cycle (resp. chain) formed by the concatenation of distinct paths in P of the same color.

A solution s of the problem consists of a set of valid chains and cycles partitioning the paths in P, expressing the particular sharing of ADMs.

More precisely, we say that two paths are adjacent if they have a common endpoint. Each path uses two ADMs, one at each endpoint; if two adjacent paths are assigned the same wavelength, then they can use the same ADM. Thus, an ADM may be shared by at most two lightpaths. In this way each valid cycle of k paths in s uses k ADMs, because every ADM is shared by exactly two paths. Similarly, each chain of k paths uses $k + 1$ ADMs, as the initial and final ADMs in the chain are used only by the initial and the final path of the chain, respectively.

We are interested in finding a solution s such that the total number of used ADMs, denoted as $ADM(s)$, is minimized.

Given an instance (G, P) and a solution s, we define the *saving graph* $S = (P, E)$, as the multigraph in which there is a node for each path $p_i \in P$ and an edge between two nodes if their corresponding paths share an ADM in s. The problem of maximizing the number of savings $SAV(s)$, i.e. the number of edges $|E|$ of the saving graph, is strictly related to the one of minimizing the total number $ADM(s)$ of used ADMs. In fact, any solution s^* maximizes $SAV(s^*)$ if and only if it minimizes $ADM(s^*)$. Moreover, the following proposition holds.

Proposition 1 *Let s^* be an optimal solution for (G, P), i.e. a solution minimizing the number of ADMs (and thus maximizing the number of savings).*

Then $\frac{ADM(s)}{ADM(s^*)} \leq 2 - r$ *for any solution* s *such that* $\frac{SAV(s)}{SAV(s^*)} = r \leq 1$.

Proof Recalling that n is the number of paths in P, we obtain

$$\frac{ADM(s)}{ADM(s^*)} = \frac{2n - SAV(s)}{2n - SAV(s^*)} = \frac{2n - r \cdot SAV(s^*)}{2n - SAV(s^*)} \leq \frac{2n - r \cdot n}{2n - n} = 2 - r,$$

where the last inequality holds because the function $\frac{2n - r \cdot SAV(s^*)}{2n - SAV(s^*)}$ is non-decreasing in $SAV(s^*)$ and $SAV(s^*) \leq n$. □

We assume that every path $p_i \in P$, $i = 1, \ldots, n$, is issued and handled by a player α_i, that for the sake of simplicity in the sequel we will often identify with p_i. At every given step a single agent α_i, by performing a selfish *move*, can decide whether and with whom to share the cost of the ADMs at the endpoints of p_i. Hence, her strategy set is the collection of all the possible subsets of at most two other adjacent (not overlapping) paths, one per endpoint. A given strategy is feasible if and only if (i) the chosen paths are not already sharing the involved ADMs with some other path and (ii) the new created chain or ring of requests induces a valid coloring, that is no two paths have an edge in common.

Clearly, a *strategy profile* $(s_1, \ldots, s_n)$ defines a solution $s \in \mathcal{S}$ of the game. A non-cooperative game $\mathcal{G}$ is defined by a tuple (G, P, f, k) where (G, P) is an optical network instance, f is a cost sharing method inducing a cost sharing function $c : \mathcal{S} \times P \rightarrow \Re$ distributing the whole hardware cost among the players and k is the maximum size of a coalition of players that can collude (notice that if $k = 1$ no collusion is allowed and thus $\mathcal{G}$ is a "classical" non-cooperative game).

We consider two fundamental cost sharing methods: the Shapley [32] ($f =$ Shapley) and the Egalitarian ($f =$ Egalitarian) ones.

In the Shapley cost sharing method, the agents sharing an ADM pay for it by equally splitting its cost. Thus, recalling that each requests needs exactly 2 ADMs, and that each ADM can be shared at most by 2 agents, the cost $c_i(s)$ charged to player α_i in the strategy profile s can be 1 (if she shares both her ADMs with other requests), $\frac{3}{2}$ (if she shares only an ADM with another request), or 2 (if she does not share any ADM with other requests). The objective of a player i is to choose the strategy minimizing her own cost, given the strategies of the other players.

In the Egalitarian cost sharing method, the whole hardware cost corresponding to a strategy profile s is divided between all the players in an egalitarian way, i.e. $c_i(s) = \frac{ADM(s)}{n}$ for every $i = 1, \ldots, n$.

Clearly in both cases, given a strategy profile s, $\sum_{i=1}^{n} c_i(s) = ADM(s)$.

If the parameter k of the game (G, P, f, k) is equal to 1, no coalition can be constituted and each player acts independently. In such a setting, a *Nash equilibrium* is a strategy profile such that no player can reduce her cost by seceding in favor of a better strategy, given the strategies of the other players. Denoting by $\mathcal{N}$ the set of all the possible Nash equilibria, the *price of anarchy* (PoA) of a game $\mathcal{G}$ is defined as the worst case ratio among the Nash versus optimal performance, i.e., $PoA(\mathcal{G}) = \frac{\max_{s \in \mathcal{N}} ADM(s)}{ADM(s^*)}$, where s^* is the strategy profile corresponding to the optimal solution. Moreover, the *optimistic price of anarchy* ($OPoA$) of $\mathcal{G}$ is defined as the best case ratio among the Nash versus the optimal performance, i.e., $OPoA(\mathcal{G}) = \frac{\min_{s \in \mathcal{N}} ADM(s)}{ADM(s^*)}$. The following proposition shows that, if $k = 1$, the games induced by the two considered cost sharing methods have the same set $\mathcal{N}$ of Nash equilibria and the same convergence behavior, and thus their prices of anarchy and convergence speeds are also equal.

Proposition 2 *Consider the games* $\mathcal{G}_1 = (G, P, \textsc{Shapley}, 1)$ *and* $\mathcal{G}_2 = (G, P, \textsc{Egalitarian}, 1)$ *defined on the same instance* (G, P). *Given a strategy profile* s, *for every* $i = 1, \ldots, n$ *player* p_i *has the same set of selfish moves to perform in* $\mathcal{G}_1$ *and* $\mathcal{G}_2$.

Proof Given a generic player p_i, we show that every selfish move of p_i starting from s in $\mathcal{G}_1$ is also a selfish move for p_i starting from s in $\mathcal{G}_2$, and vice versa.

Consider the saving graph $S = (P, E)$ relative to s and let δ_i the degree of node p_i in S. Clearly, the cost charged to p_i according to the Shapley cost sharing method is $2 - \frac{\delta_i}{2}$.

If p_i has a selfish move M according to the cost sharing method Shapley, then, by performing M, increases δ_i by $\Delta > 0$. Thus, since the number of edges in the saving graph increases after p_i performs M, M is a selfish move according also to the cost sharing method Egalitarian.

On the other hand, if p_i has a selfish move M according to the cost sharing method Egalitarian, she, by performing M, increases by $\Delta > 0$ the number of edges of S. The only way to obtain such an growth is by increasing her degree δ_i. Thus, M is a selfish move according also to the cost sharing method Shapley. □

For the sake of clearness, we will assume in all the proofs related to the price of anarchy determination that the cost sharing method is the Shapley one.

If the parameter k of the game (G, P, f, k) is greater than 1, a *Nash equilibrium* is a strategy profile such that no coalition of k player can reduce its whole cost (sum of single costs) by seceding in favor of a better strategy, given the strategies of the other $n-k$ players. In such a setting, denoting by $\mathcal{N}_k$ the set of all the possible Nash equilibria with coalitions of size at most k, the *price of collusion* (PoC) of a game $\mathcal{G}$ is defined as the worst case ratio among the Nash versus optimal performance, i.e., $PoC_k(\mathcal{G}) = \frac{\max_{s \in \mathcal{N}_k} ADM(s)}{ADM(s^*)}$, where s^* is the strategy profile corresponding to the optimal solution. Moreover, the *optimistic price of collusion* $(OPoC)$ of $\mathcal{G}$ is defined as the best case ratio among the Nash versus the optimal performance, i.e., $OPoC_k(\mathcal{G}) = \frac{\min_{s \in \mathcal{N}_k} ADM(s)}{ADM(s^*)}$. Notice that, since the coalitions can dynamically change, in this case the Shapley cost sharing method is not well defined. Thus, for $k > 1$ we will focus only on the games induced by the egalitarian cost sharing method.

Let us now present some preliminary results about the existence and convergence to Nash Equilibria. In particular, we show that every game always converges to a Nash equilibrium in a linear number of moves.

Proposition 3 *In every game $\mathcal{G} = (G, P, f, k)$, where $f \in \{$*SHAPLEY, EGALITARIAN*$\}$ and $k = 1$ or $f =$* EGALITARIAN *and $k \geq 2$, the social function ADM is a potential function, i.e. if s' is the strategy profile resulting from the strategy profile s after the selfish moves of the colluding players $\alpha_{i_1}, \ldots, \alpha_{i_h}$, with $h \leq k$, $ADM(s') < ADM(s)$.*

Proof Notice that by Proposition 2, for the case $k = 1$ it is sufficient to prove the claim for the egalitarian cost sharing method.

By the selfishness of the moves of the colluding players, we have that

$$c_{i_1}(s') = \cdots = c_{i_h}(s') = \frac{ADM(s')}{n} < \frac{ADM(s)}{n} = c_{i_1}(s) = \cdots = c_{i_h}(s),$$

and the claim easily follows. □

As a direct consequence of the previous propositions, by also considering Proposition 2, since the social function, that is a potential function for the game, can assume at most $n + 1$ different values, it holds that every game always converges to a Nash equilibrium in at most n selfish moves.

Moreover, since the optimal solution is a minimum of the defined potential functions, it follows that the optimal solution is also an equilibrium. Therefore, given any game $\mathcal{G}$, the *optimistic* price of anarchy is the best possible one, i.e. $OPoA(\mathcal{G}) = 1$, and, for every integer $k > 1$, the same holds for the *optimistic* price of collusion, i.e. $OPoC_k(\mathcal{G}) = 1$. Therefore, in the remaining part of the paper we will focus on the price of anarchy (for $k = 1$) and price of collusion (for $k > 1$).

3 Price of Anarchy

In this section we analyze the price of anarchy for the ADM minimization problem on games in which no collusion between players is allowed, and thus each selfish player acts independently. More precisely we prove that the price of anarchy is at most $\frac{5}{3}$, and that such a bound is tight even for ring network topologies.

Lemma 4 *Given a game* $\mathcal{G} = (G, P, f, 1)$ *with* $f \in \{\text{SHAPLEY}, \text{EGALITARIAN}\}$ *and a Nash Equilibrium* s *for* $\mathcal{G}$, $\frac{SAV(s)}{SAV(s^*)} \geq \frac{1}{3}$, *where* s^* *is the strategy profile corresponding to an optimal solution.*

Proof Consider the saving graph $S^* = (P, E^*)$ corresponding to the strategy profile s^*. Moreover, let $S = (P, E)$ be the saving graph corresponding to the strategy profile s. For every $i = 1, \ldots, n$, let δ_i^* and δ_i be the degree of node p_i in S^* and S, respectively; we want to prove that

$$\frac{\sum_{i=1}^{n} \delta_i}{\sum_{i=1}^{n} \delta_i^*} = \frac{2SAV(s)}{2SAV(s^*)} \geq \frac{1}{3}.$$

We now show two interesting properties relating S and S^*.

(1) If $(p_a, p_b) \in E^*$, then $\delta_a + \delta_b \geq 1$.
In fact, otherwise s would not be a Nash Equilibrium because p_a and p_b would have incentive to share an ADM.

(2) If $(p_a, p_b) \in E^*$ and $(p_b, p_c) \in E^*$, then $\delta_a + \delta_b + \delta_c \geq 2$.
If there exists $i \in \{a, b, c\}$ such that $\delta_i = 2$, or there exist $i, j \in \{a, b, c\}, i \neq j$ such that $\delta_i = \delta_j = 1$, the property holds. Moreover, since by the previous property $\delta_a + \delta_b \geq 1$ and $\delta_b + \delta_c \geq 1$ we obtain $\delta_a + 2\delta_b + \delta_c \geq 2$; thus, if $\delta_b = 0$ we also obtain the desired property. It remains to analyze the case in which $\delta_b = 1$ and $\delta_a = \delta_c = 0$. Clearly, this case is not possible because s would not be a Nash Equilibrium; in fact, p_b would have incentive to leave her current strategy and share an ADM with p_a and another ADM with p_c.

Since each node of a saving graph can have degree at most two, we can partition E^* into h cycles $E_1^*, \ldots, E_h^*$ and l paths $E_1'^*, \ldots, E_l'^*$, and consequently P in $P_0, P_1, \ldots, P_h, P_1', \ldots, P_l'$, such that P_0 contains all the nodes having degree 0 in S^*, P_j, for every $j = 1, \ldots, h$, contains all the nodes corresponding to E_j^*, i.e. all the nodes having at least an edge belonging to E_j^* incident to them, and P_j', for every $j = 1, \ldots, l$, contains all the nodes corresponding to $E_j'^*$. In the remaining part of the proof, we show that for each set $X \in \{P_1, \ldots, P_h, P_1', \ldots, P_l'\}$, $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{1}{3}$, thus proving the claim. The proof is divided in several distinct cases.

- $|X| = 3m$, with $m = 1, 2, \ldots$; $X \in \{P_1, \ldots, P_h, P_1', \ldots, P_l'\}$
  By exploiting Property 2, since in the corresponding component of the optimal solution the sum of the degrees is at most $6m$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{2m}{6m} = \frac{1}{3}$.
- $|X| = 3m + 1$, with $m = 1, 2, \ldots$; $X \in \{P_1', \ldots, P_l'\}$
  By exploiting Property 2, since in the corresponding path of the optimal solution the sum of the degrees is exactly $6m$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{2m}{6m} = \frac{1}{3}$.
- $|X| = 3m + 2$, with $m = 0, 1, \ldots$; $X \in \{P_1', \ldots, P_l'\}$
  By exploiting Property 2 for the first $3m$ nodes and Property 1 for the last 2 nodes, since in the corresponding path of the optimal solution the sum of the degrees is exactly $6m + 2$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{2m+1}{6m+2} > \frac{1}{3}$.
- $|X| = 3m + 1$, with $m = 1, 2, \ldots$; $X \in \{P_1, \ldots, P_h\}$
  Clearly, by Property 1 at least a node of the cycle must have degree at least 1. By exploiting Property 2 for the remaining m nodes of the cycle, since in the corresponding cycle of the optimal solution the sum of the degrees is exactly $6m + 2$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{2m+1}{6m+2} = \frac{1}{3}$.
- $|X| = 3m + 2$, with $m = 0, 1, \ldots$; $X \in \{P_1, \ldots, P_h\}$
  First of all, we show that there must exist two consecutive nodes p_a and p_b of the considered cycle such that $\delta_a + \delta_b \geq 2$. If there exists a node having degree 2, this property is trivially verified. Otherwise, all nodes have degree at most 1, and we want to show that it is not possible that two consecutive nodes having degree 1 do not exist. If $m = 0$, it is easy to check that, since s is an equilibrium, the sum of the degrees of the two nodes in X is at least 2. If $m > 0$, let p_a be a node of the cycle having degree 1 (by Property 1 it necessarily exists). If both the two adjacent nodes in the cycle of the optimal solution have degree 0 in S, s would not be an equilibrium since p_a would change her current strategy by sharing both her ADMs with such adjacent nodes. Therefore, letting p_b be the adjacent node with degree $\delta_b > 0$, we obtain $\delta_a + \delta_b \geq 2$.
  By exploiting Property 2 for the remaining m nodes of the cycle, since

in the corresponding cycle of the optimal solution the sum of the degrees is exactly $6m+4$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{2m+2}{6m+4} > \frac{1}{3}$.

□

By combining the previous lemma with Proposition 1, the following theorem holds.

Theorem 5 *For any game $\mathcal{G} = (G, P, f, 1)$, with $f \in$ {*Shapley*, *Egalitarian*}, $PoA(\mathcal{G}) \leq \frac{5}{3}$.*

Now we provide a matching lower bound, holding for a network having ring topology. The following theorem proves that the previous upper bound is tight even for ring networks, and thus the price of anarchy is equal to $\frac{5}{3}$.

Theorem 6 *For any $\epsilon > 0$, there exists an instance of the ADM minimization game $\mathcal{G} = (G, P, f, 1)$, where $f \in$ {*Shapley*, *Egalitarian*} and G is a ring network, such that $PoA(\mathcal{G}) \geq \frac{5}{3} - \epsilon$.*

Proof Given two generic requests α and β, we denote by $C(\alpha, \beta)$ the set of their common endpoints and by $I(\alpha, \beta)$ the set of edges in which they are overlapping.

In order to prove the theorem, we provide an inductive construction in which at each step $i = 0, 1, \ldots$ three requests a_i,b_i and c_i forming a cycle are added (for $i = 0$ they do not form a cycle). We want to show that at each step i, there exists a solution s_i^* such that $ADM(s_i^*) = 5+3i$, and there exists a Nash equilibrium s_i such that $ADM(s_i) = 5 + 5i$; by letting i go to the infinity, the price of anarchy tends to $\frac{5}{3}$.

Let a_0,b_0 and c_0 be such that $C(a_0, b_0) \neq \emptyset$, $C(a_0, c_0) \neq \emptyset$ and $I(c_0, b_0) \neq \emptyset$. It is easy to check that there exist a solution s_0^* using 5 ADMs and being also a Nash equilibrium; thus, $s_0 = s_0^*$ and $ADM(s_0) = ADM(s_0^*) = 5$. Moreover, let a_1,b_1 and c_1 be 3 requests forming a cycle, such that $C(a_1, b_1) = C(a_0, c_0)$, $I(a_1, a_0) \neq \emptyset$, $b_1 \cap I(b_0, c_0) \neq \emptyset$, $I(c_1, \beta) \neq \emptyset$ for $\beta \in \{a_0, b_0, c_0\}$ and $I(b_1, \beta) \neq \emptyset$ for $\beta \in \{b_0, c_0\}$. Clearly, there exists a solution s_1^* using 8 ADMs. On the other hand, the configuration s_1 in which a_1 and c_0 share an ADM, b_1 and a_0 share another ADM and b_0 and c_1 do not share any ADM is a Nash equilibrium of cost 10.

Now we prove by induction that at each step $i \geq 2$ we can add three requests a_i,b_i and c_i forming a cycle, such that the cost of the optimal solution is at most $5+3i$, and there exists a Nash equilibrium s_i of cost $ADM(s_i) = 5+5i$. More precisely, we show by induction that for every $i \geq 2$ there exist a_i,b_i and c_i such that: (i) $C(a_i, b_i) = C(a_{i-1}, c_{i-1})$; (ii) $I(a_i, a_{i-1} \neq \emptyset)$; (iii) $b_i \cap I(c_{i-1}, c_{i-2}) \neq \emptyset$; (iv) $I(c_i, \beta) \neq \emptyset$ for $\beta \in \{a_{i-1}, b_{i-1}, c_{i-1}\}$; (v) $I(b_i, \beta) \neq \emptyset$ for $\beta \in \{c_{i-2}, c_{i-1}\}$; (vi) there exists a Nash equilibrium s_i of cost $ADM(s_i) = 2(i+1)+3(i+1) = 5+5i$ and there exists a solution s_i^* of cost $ADM(s_i^*) = 5+3i$.

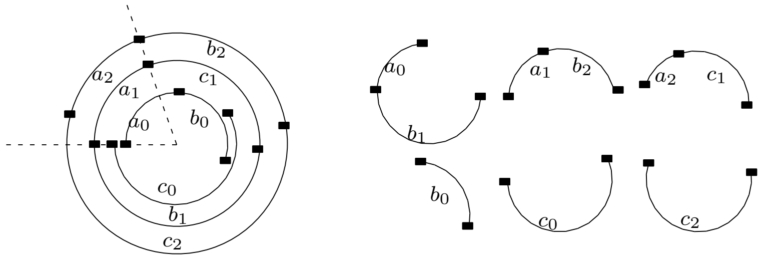

(a) (b)

Fig. 1. (a) Solution s_2^* (b) Solution s_2

For $i = 2$, the base of the induction is verified. In fact, (i) and (ii) can trivially hold by construction; since $I(c_1, c_0) \neq \emptyset$, b_2 can be chosen such that also (iii) holds; (iv) follows from (ii) and (iii) as by (ii) $I(c_2, a_1) \neq \emptyset$ and by (iii) $I(c_2, b_1) \neq \emptyset$ and $I(c_2, c_1) \neq \emptyset$; (v) follows from (iii). Moreover, it can be checked that the configuration s_2 in which the couples of requests a_2 and c_1, b_2 and a_1, a_0 and b_1 share an ADM, and requests b_0, c_0 and c_2 share no ADM is a Nash equilibrium of cost 15 and there exists a configuration s_2^*, in which the requests added at step 1 and the ones added at step 2 form two cycles, such that $ADM(s_2^*) = 11$ (see Figure 1).

Assuming the inductive claim true for any $j \leq i$, we now prove it for $i + 1$.

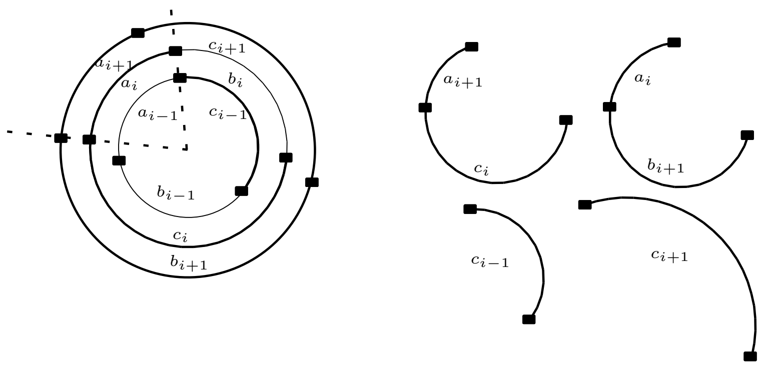

(a) (b)

Fig. 2. (a) Solution s_{i+1}^* (b) Solution s_{i+1}

Clearly, a_{i+1}, b_{i+1} and c_{i+1} can be chosen such that conditions (i) and (ii) hold (see Figure 2); moreover, since by the condition (iv) of the inductive hypothesis concerning step i $I(c_i, c_{i-1}) \neq \emptyset$, b_{i+1} can be chosen such that $b_{i+1} \cap I(c_i, c_{i-1}) \neq \emptyset$, and thus also (iii) holds. By (ii), since $C(a_{i+1}, c_{i+1}) \neq \emptyset$,

$I(c_{i+1}, a_i) \neq \emptyset$ and by (iii) $I(c_{i+1}, b_i) \neq \emptyset$ and $I(c_{i+1}, c_i) \neq \emptyset$; thus (iv) holds. (v) directly follows from (iii).

In order to prove (vi), let us first notice that since at each step having index greater than 0 three requests forming a cycle are added, we obtain that there exists a solution s^*_{i+1} such that $ADM(s^*_{i+1}) = 5 + 3(i + 1)$.

For every $j = 1, \ldots, i + 1$, it follows from conditions (i), (ii) and (iii) of the inductive hypothesis concerning step j that $I(b_j, a_{j-1}) = \emptyset$; moreover, exploiting a similar argument, it holds that $I(a_{i+1}, c_i) = \emptyset$. Consider the strategy profile s_{i+1} in which each couple of requests (b_j, a_{j-1}) for $j = 1, \ldots, i + 1$ and the couple of request (a_{i+1}, c_i) share an ADM; moreover, in s_{i+1} requests $b_0, c_0, \ldots, c_{i-1}, c_{i+1}$ do not share any ADM. Clearly, $ADM(s_{i+1}) = 5+5(i+1)$; it remains to prove that s_{i+1} is an equilibrium.

Since in s_{i+1} there exists no couple of requests added to the instance at the same step, all the coupled requests are in equilibrium as none of them can share two ADMs by changing her strategy. Moreover, b_0 cannot change her strategy by sharing an ADM with a_0 because a_0 is coupled with b_1, and b_1 and b_0 are overlapping requests ($I(b_0, b_1) \neq \emptyset$). For every $j = 0, \ldots, i-1$, c_j cannot change her strategy by sharing an ADM neither with a_j nor with b_j: c_j cannot share an ADM with a_j because a_j is coupled with b_{j+1} and by condition (v) of the inductive hypothesis concerning step $j + 1$, $I(b_{j+1}, c_j) \neq \emptyset$; c_j cannot share an ADM with b_j because b_j is coupled with a_{j-1} and by condition (iv) of the inductive hypothesis concerning step j, $I(c_j, a_{j-1}) \neq \emptyset$. Finally, c_{i+1} cannot change her strategy by sharing an ADM neither with a_{i+1} nor with b_{i+1} because a_{i+1} is coupled with c_i and b_{i+1} is coupled with a_i and by condition (iv), $I(c_{j+1}, c_i) \neq \emptyset$ and $I(c_{j+1}, a_i) \neq \emptyset$. □

4 Price of Collusion

In this section we analyze the price of collusion for the ADM minimization problem on games in which coalitions of at most k players can collude. More precisely we prove that the price of collusion is between $\frac{3}{2}$ and $\frac{3}{2} + \frac{1}{k}$, with $\frac{3}{2} + \epsilon$ being the approximation guaranteed by the best know approximation algorithms [8,17] (polynomial for every $\epsilon > 0$) for this problem on general network topologies. As already remarked, the evolution of such games naturally defines, for every fixed k, a polynomial time approximation algorithm for the ADM minimization problem. Thus, such games interestingly define a distributed algorithm with the same approximation guaranteed of the best know centralized algorithm.

In order to provide an upper bound to the price of collusion, we exploit arguments similar to the ones used in the proof of Lemma 4.

Lemma 7 *For every $k = 2, 3, \dots$, given a game $\mathcal{G} = (G, P, \text{EGALITARIAN}, k)$ and a Nash Equilibrium s for $\mathcal{G}$, $\frac{SAV(s)}{SAV(s^*)} \geq \frac{1}{2} - \frac{1}{k}$, where s^* is the strategy profile corresponding to an optimal solution.*

Proof Consider the saving graph $S^* = (P, E^*)$ corresponding to the strategy profile s^*. Moreover, let $S = (P, E)$ be the saving graph corresponding to the strategy profile s. For every $i = 1, \dots, n$, let δ_i^* and δ_i be the degree of node p_i in S^* and S, respectively; we want to prove that

$$\frac{\sum_{i=1}^n \delta_i}{\sum_{i=1}^n \delta_i^*} = \frac{2SAV(s)}{2SAV(s^*)} \geq \frac{1}{2} - \frac{1}{k}.$$

A key property relating S and S^* holds: for any $r \leq k$, if for every $i = 1, \dots, r-1$ $(p_{a_i}, p_{a_{i+1}}) \in E^*$, then $\sum_{i=1}^r \delta_{a_i} \geq r - 1$. In fact, otherwise $\sum_{i=1}^r \delta_{a_i} \leq r - 2$ and by changing their strategy players $p_{a_1}, \dots, p_{a_r}$ can remove at most $r - 2$ edges from E; moreover, they can change their strategy so as to add for every $i = 1, \dots, r-1$ edge $(p_{a_i}, p_{a_{i+1}})$ to E. Thus, by recalling that every player in s is charged with a cost equal to $\frac{ADM(s)}{n} = \frac{2n - SAV(s)}{n}$, since players $p_{a_1}, \dots, p_{a_r}$ by colluding could evolve in a strategy profile s' with one more edge, i.e. with one more saving, s would not be an equilibrium.

Since each node of a saving graph can have degree at most two, we can partition E^* into h cycles $E_1^*, \dots, E_h^*$ and l paths $E_1'^*, \dots, E_l'^*$, and consequently P in $P_0, P_1, \dots, P_h, P_1', \dots, P_l'$, such that P_0 contains all the nodes having degree 0 in S^*, P_j, for every $j = 1, \dots, h$, contains all the nodes corresponding to E_j^*, i.e. all the nodes having at least an edge belonging to E_j^* incident to them, and P_j', for every $j = 1, \dots, l$, contains all the nodes corresponding to $E_j'^*$. In the remaining part of the proof, we show that for each set $X \in \{P_1, \dots, P_h, P_1', \dots, P_l'\}$, $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{1}{2} - \frac{1}{k}$, thus proving the claim. Let $|X| = m \cdot k + q$, with m and q non-negative integers; the proof is divided in three distinct cases.

- $m \geq 1$ and $q = 0$
  We can partition the nodes in X in m sets of size k, and by exploiting the key property, since in the corresponding component of the optimal solution the sum of degrees is at most $2mk$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{m(k-1)}{2mk} > \frac{1}{2} - \frac{1}{k}$.
- $m \geq 1$ and $q \geq 1$
  We can partition the nodes in X in m sets of size k and a set of size q, and by exploiting the key property, since in the corresponding component of the optimal solution the sum of degrees is at most $2(mk + q)$, we have

that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{m(k-1)+q-1}{2(mk+q)} = \frac{1}{2} - \frac{m+1}{2(mk+q)} \geq \frac{1}{2} - \frac{1}{k}$.

- $m = 0$

  In this case, $q = |X| < k$ and the key property can be strengthened as follows: $\sum_{p_i \in X} \delta_i \geq q$. In order to prove this strengthened property, it is sufficient to apply an argument very similar to the one exploited for the key property, by noticing that in this case if all players in X collude, they can add to E q edges (instead of $q-1$).

  Thus, since in the corresponding component of the optimal solution the sum of degrees is at most $2q$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{q}{2q} = \frac{1}{2} > \frac{1}{2} - \frac{1}{k}$.

□

By combining the previous lemma with Proposition 1, the following theorem holds.

Theorem 8 *For every $k = 2, 3, \ldots$, any game $\mathcal{G} = (G, P, \text{EGALITARIAN}, k)$ is such that $PoC_k(\mathcal{G}) \leq \frac{3}{2} + \frac{1}{k}$.*

The following theorem provides an almost matching lower bound.

Theorem 9 *For every $k = 2, 3, \ldots$, there exists an instance of the ADM minimization game $\mathcal{G} = (G, P, \text{EGALITARIAN}, k)$ such that $PoC_k(\mathcal{G}) \geq \frac{3}{2}$.*

Proof We first introduce some definitions and preliminary results needed in the following. Let us define the labelled undirected *shareability graph* $G_S = (P, E_S)$ of an instance as the union of the saving graphs of all possible solutions of this instance. In other words, there is an edge in the shareability graph if and only if the corresponding two paths can be connected to share an ADM in some node v of G, and this edge is labelled v. Obviously, given some node p of G_S, the label set of its incident edges has size at most two.

The *conflict graph* $G_C = (P, E_C)$ of an instance has an edge between two nodes p and p' if and only if the paths p and p' overlap in G. Note that $E_S \cap E_C = \emptyset$, because no solution can share an ADM between p and p' if they overlap. For the correspondence between instances of the ADM minimization problem and the shareability and conflict graphs see [17].

A (k, g)-*cage* is a regular graph of degree k and girth g with minimal number of vertices, and $v(k, g)$ is its number of vertices. In [30] is shown that $v(k, g)$ is finite, and an upper bound is given in [15].

Instead of describing the instance claimed, we will describe its shareability and conflict graphs.

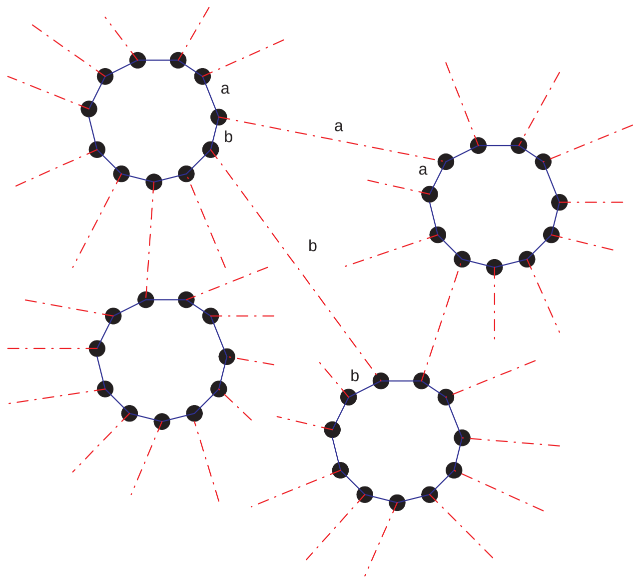

Fig. 3. The Graph G_S constructed from some cage H

For any integer $k \geq 1$, we choose H be a $(k+1, k+1)$-cage with $n = v(k+1, k+1)$ nodes. For each node in H there is a ring of $k+1$ nodes and $k+1$ *blue* edges in G_S. Obviously, each node has exactly two blue edges incident to it. For each edge (u, v) of H we add a *red* edge to G_S as follows. Choose an arbitrary node u' (resp. v') from the cycle corresponding to u (resp. v) which does not have an incident red edge. Add the red edge (u', v') to G_S. Because H is $(k+1)$-regular, each node of G_S has exactly one red edge incident to it. G_S is the graph consisting of the red and blue edges. We label the red edges with distinct labels. In order to label the blue edges we assume a direction for each cycle. A blue edge is labelled by the label of the red edge incident to the node in the clockwise direction. See Figure 3.

G_C contains all the possible edges between nodes from different rings, except the (red) edges of G_S.

The solution s^* consisting of the blue edges of G_S is feasible, because the blue rings are independent sets of of G_C, and the labels of two consecutive blue edges are distinct. Moreover, it is optimal, because its sharing graph consists of the maximum number of edges possible, namely $SAV(s^*) = n(k+1)$. Therefore $ADM(s^*) = 2n(k+1) - n(k+1) = n(k+1)$.

On the other hand the solution s consisting of the red edges of G_S is feasible because each connected component is an edge with two nodes, and the paths corresponding to these two nodes are not overlapping. The number of edges in this solution is $SAV(s) = \frac{n(k+1)}{2}$. Therefore $ADM(s) = 2n(k+1) - \frac{n(k+1)}{2} = \frac{3}{2}n(k+1)$ and $\frac{ADM(s)}{ADM(s^*)} = \frac{3}{2}$. It remains to prove that s is a Nash Equilibrium.

Assume, by contradiction that there is some coalition of $k' \leq k$ nodes, that can improve its gain by choosing some other strategy profile implying a solution $s' \neq s$. Let a *segment* be a set of consecutive nodes in some blue ring. Let l be the number of maximal segments in the coalition. Obviously $1 \leq l \leq k'$. For any solution s' and $1 \leq i \leq l$, let $e_i(s')$ be the number of blue edges of s' connecting nodes of segment i. For solution s we have $\forall i, e_i(s) = 0$. Let $l' \leq l$ be the number of segments such that $e_i(s') > 0$.

Note that, because of the way we constructed the conflict graph, a red edge incident to some node v, chosen in solution s' implies that the two blue edges incident to v are not in s'. A blue ring has size $k+1$, thus a segment may not span an entire blue ring. Therefore when the coalition changes its strategy profile so that the solution changes from s to s', the number of red edges incident to segment s_i decreases by $\Delta_{i,red} \geq e_i(s') + 1$. On the other hand the number of blue edges increases by $\Delta_{i,blue} = e_i(s')$. Thus the number of edges incident to such a segment decreases by $\Delta_i = \Delta_{i,red} - \Delta_{i,blue} \geq 1$. Therefore the total number of edges incident to all segments decreases by $\Delta \geq l' - x$, where x is the number of red edges connecting these l' segments. This is because in $\sum \Delta_i \geq l'$, these x edges might be counted twice.

Let $l'' \leq l'$ be the number of blue rings containing the l' segments. Assume $l'' > 0$. Consider the edges in H corresponding to the red edges connecting these rings. They form a subgraph of H with l'' nodes. But $l'' \leq l' \leq l \leq k' \leq k < k+1$, and H has girth $k+1$, thus this subgraph is acyclic and contains at most $l''-1$ edges. Then, $x \leq l''-1 \leq l'-1$, and $\Delta \geq l' - x \geq l' - (l'-1) = 1$. Namely the coalition's gain decreases, a contradiction. Therefore $l'' = 0$. The number of blue rings containing the l' segments is zero, therefore the number of these segments $l' = 0$. In other words, $\forall i, e_i(s') = 0$. Namely the edge set of s', incident to the coalition, consists of red edges only. But s already contains all the red edges incident to the nodes of the coalition, then the gain of the coalition may not increase: a contradiction. □

5 Local Search and Concluding Remarks

In this section we show some basic results emphasizing that local search is a promising approach for possibly improving the achievable approximation ratio of the ADMs minimization problem.

As already remarked, under the assumption of collusion of at most k players, the Shapley method does not induce well defined games. This stems on the fact that the payment of a player is not solely a function of the current strategy profile, but is also affected by the history of the past coalitions. However, Shapley naturally yields local search schema with the induced definition of neighborhood of a current solution s. Namely, any solution s' that can be obtained from s by modifying the strategy of at most k players is a neighbor of s; such a solution is an improving one with respect to s and the fixed coalition if it reduces the sum of the Shapley costs of the involved players, that is it increases the sum of their degrees in the saving graph.

The following proposition characterizes the performance of local optima according to such a neighborhood definition.

Proposition 10 *For every $k = 2, 3, \ldots,$ any local optimum solution s in the schema induced by the above definition of neighborhood has total cost $ADM(s) \leq \left(1 + \frac{2}{k}\right) ADM(s^*)$, where s^* is an optimal solution.*

Proof By Proposition 1, it suffices to show that $\frac{SAV(s)}{SAV(s^*)} \geq 1-\frac{2}{k}$. Consider the saving graph $S^* = (P, E^*)$ corresponding to the optimal solution s^*. Moreover, let $S = (P, E)$ be the saving graph corresponding to s. For every $i = 1, \ldots, n$, let δ_i^* and δ_i be the degree of node p_i in S^* and S, respectively; we want to prove that

$$\frac{\sum_{i=1}^n \delta_i}{\sum_{i=1}^n \delta_i^*} = \frac{2SAV(s)}{2SAV(s^*)} \geq 1 - \frac{2}{k}.$$

A key property relating S and S^* holds: for any $r \leq k$, if for every $i = 1, \ldots, r-1$ $(p_{a_i}, p_{a_{i+1}}) \in E^*$, then $\sum_{i=1}^r \delta_{a_i} \geq 2(r-1)$. In fact, otherwise players $p_{a_1}, \ldots, p_{a_r}$ can change their strategy so as to add for every $i = 1, \ldots, r-1$ edge $(p_{a_i}, p_{a_{i+1}})$ to E, thus obtaining sum of degrees equal to $2(r-1)$.

Since each node of a saving graph can have degree at most two, we can partition E^* into h cycles $E_1^*, \ldots, E_h^*$ and l paths $E_1'^*, \ldots, E_l'^*$, and consequently P in $P_0, P_1, \ldots, P_h, P_1', \ldots, P_l'$, such that P_0 contains all the nodes having degree 0 in S^*, P_j, for every $j = 1, \ldots, h$, contains all the nodes corresponding to E_j^*, i.e. all the nodes having at least an edge belonging to E_j^* incident to them, and P_j', for every $j = 1, \ldots, l$, contains all the nodes corresponding to $E_j'^*$. In the remaining part of the proof, we show that for each set $X \in \{P_1, \ldots, P_h, P_1', \ldots, P_l'\}$, $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq 1 - \frac{2}{k}$, thus proving the claim. Let $|X| = m \cdot k + q$, with m and q non-negative integers; the proof is divided in three distinct cases.

- $m \geq 1$ and $q = 0$
  We can partition the nodes in X in m sets of size k, and by exploiting the key property, since in the corresponding component of the optimal solution the sum of degrees is at most $2mk$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{2m(k-1)}{2mk} = 1 - \frac{1}{k} > 1 - \frac{2}{k}$.
- $m \geq 1$ and $q \geq 1$

We can partition the nodes in X in m sets of size k and a set of size q, and by exploiting the key property, since in the corresponding component of the optimal solution the sum of degrees is at most $2(mk+q)$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{2m(k-1)+2(q-1)}{2(mk+q)} = 1 - \frac{m+1}{mk+q} \geq 1 - \frac{2}{k}$.

- $m = 0$

  In this case, $q = |X| < k$ and the key property can be strengthened as follows: $\sum_{p_i \in X} \delta_i \geq 2q$. In order to prove this strengthened property, it is sufficient to apply an argument very similar to the one exploited for the key property, by noticing that in this case if all players in X collude, they can add to E q edges (instead of $q-1$).

  Thus, since in the corresponding component of the optimal solution the sum of degrees is at most $2q$, we have that $\frac{\sum_{p_i \in X} \delta_i}{\sum_{p_i \in X} \delta_i^*} \geq \frac{2q}{2q} = 1 > 1 - \frac{2}{k}$.

□

Unfortunately, such a neighbor definition for increasing values of k does not induce a PTAS, since the schema not only does not converge in a polynomial number of steps, might not converge at all and local optima may even not exist.

Proposition 11 *The local search schema induced by the above definition of neighborhood may posses no local optimum.*

Proof Let G be a ring network with two nodes, and consider 3 requests p_1, p_2, p_3 between such nodes, such that p_2 and p_3 collide and p_1 can connect either to p_2 or to p_3.

We now show that no solution is a local optimum.

If p_1 does not connect to any other request, there exists an improving step in which p_1 and p_2 (or p_3) connect each to the other.

If p_1 is connected to p_2 (resp. p_3), there exists an improving step in which p_1 and p_3 (resp. p_2) connect each to the other. □

The above results on local search emphasize that the determination of new cost sharing methods reaching a compromise between the Shapley and Egalitarian ones in terms of optimization and performance is a promising and worth investigating issue. To this aim we observe that a linear combination of the two criteria is affected by the same unconvergence behavior. In fact, in the instance shown in the proof of Proposition 11, the solutions in which two requests connect can be involved in a cycle of improving steps and have the same total cost: since the Egalitarian contribution in the linear combination is fixed, the Shapley part causes exactly the same behavior. Nevertheless, the determination of other intermediate methods combining both the Shapley and Egalitarian advantages is an important left open question.

Besides the above mentioned results for general topologies, it would be also nice to determine specific collusion results for ring networks, possibly improving the related approximation ratios.

Finally, a last interesting issue is that of extending our results to the grooming case in which up to a certain number of paths g of the same color can share the same physical links and the same ADMs [12].

References

[1] Elliot Anshelevich, Anirban Dasgupta, Eva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. In *STOC '03: Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 511–520, New York, NY, USA, 2003. ACM Press.

[2] R. Aumann. Acceptable points in general cooperative n-person games. *Contributions to the Theory of Games, Princeton University Press*, IV, 1959.

[3] B. Beauquier, J.-C. Bermond, L. Gargano, P. Hell, S. Perennes, and U. Vaccaro. Graph problems arising from wavelength–routing in all–optical networks. In *Proc. 2nd Workshop on Optics and Computer Science, WOCS'97*, April 1997.

[4] V. Bilò and L. Moscardelli. The price of anarchy in all-optical networks. In *Proceedings of the 11th Colloquium on Structural Information and Communication Complexity (SIROCCO)*, volume 3104 of *LNCS*, pages 13–22. Springer, 2004.

[5] Vittorio Bilò, Michele Flammini, and Luca Moscardelli. On nash equilibria in non-cooperative all-optical networks. In Volker Diekert and Bruno Durand, editors, *STACS*, volume 3404 of *Lecture Notes in Computer Science*, pages 448–459. Springer, 2005.

[6] C. A. Brackett. Dense wavelength division multiplexing networks: principles and applications. *IEEE Journal on Selected Areas in Communications*, 8:948–964, 1990.

[7] N. K. Chung, K. Nosu, and G. Winzer. Special issue on dense wdm networks. *IEEE Journal on Selected Areas in Communications*, 8, 1990.

[8] G. Călinescu, Ophir Frieder, and Peng-Jun Wan. Minimizing electronic line terminals for automatic ring protection in general wdm optical networks. *IEEE Journal of Selected Area on Communications*, 20(1):183–189, Jan 2002.

[9] G. Călinescu and P-J. Wan. Traffic partition in wdm/sonet rings to minimize sonet adms. *Journal of Combinatorial Optimization*, 6(4):425–453, 2002.

[10] A. Czumaj and B. Vocking. Tight bound for the worst-case equilibria. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Mathematics (SODA)*, pages 413–420. SIAM, 2002.

[11] D. H. C. Du and R. J. Vetter. Distributed computing with high-speed optical networks. In *Proceeding of IEEE Computer*, volume 26, pages 8–18, 1993.

[12] R. Dutta and G. N. Rouskas. Traffic grooming in wdm networks: past and future. *IEEE Network*, pages 44–56, Nov.-Dec. 2002.

[13] T. Eilam, S. Moran, and S. Zaks. Lightpath arrangement in survivable rings to minimize the switching cost. *IEEE Journal of Selected Area on Communications*, 20(1):172–182, Jan 2002.

[14] L. Epstein and A. Levin. Better bounds for minimizing sonet adms. In *2nd Workshop on Approximation and Online Algorithms, Bergen, Norway*, September 2004.

[15] P. Erdös and H. Sachs. Reguláre graphen gegebener taillenweite mit minimaler knotenzahl. *Wiss. Z. Uni. Halle (Math. Nat.)*, 12:251–257, 1963.

[16] A. Fanelli, M. Flammini, G. Melideo, and L. Moscardelli. Game theoretical issues in optical networks. In *Proceedings of the 8th International Conference on Transparent Optical Networks (ICTON), Nottingham, United Kingdom.* IEEE, 2006.

[17] Michele Flammini, Mordechai Shalom, and Shmuel Zaks. On minimizing the number of adms in a general topology optical network. In Shlomi Dolev, editor, *DISC*, volume 4167 of *Lecture Notes in Computer Science*, pages 459–473. Springer, 2006.

[18] Dimitris Fotakis, Spyros C. Kontogiannis, and Paul G. Spirakis. Atomic congestion games among coalitions. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP (1)*, volume 4051 of *Lecture Notes in Computer Science*, pages 572–583. Springer, 2006.

[19] L. Gargano and U. Vaccaro. *"Routing in All–Optical Networks: Algorithmic and Graph–Theoretic Problems"in: Numbers, Information and Complexity.* Kluwer Academic, 2000.

[20] George F. Georgakopoulos, Dimitris J. Kavvadias, and Leonidas G. Sioutis. Nash equilibria in all-optical networks. In Xiaotie Deng and Yinyu Ye, editors, *WINE*, volume 3828 of *Lecture Notes in Computer Science*, pages 1033–1045. Springer, 2005.

[21] O. Gerstel, P. Lin, and G. Sasaki. Wavelength assignment in a wdm ring to minimize cost of embedded sonet rings. In *INFOCOM'98, Seventeenth Annual Joint Conference of the IEEE Computer and Communications Societies*, 1998.

[22] P. E. Green. *Fiber-Optic Communication Networks.* Prentice Hall, 1992.

[23] Ara Hayrapetyan, Éva Tardos, and Tom Wexler. The effect of collusion in congestion games. In Jon M. Kleinberg, editor, *STOC*, pages 89–98. ACM, 2006.

[24] R. Klasing. Methods and problems of wavelength-routing in all-optical networks. In *Proceeding of the MFCS'98 Workshop on Communication, August 24-25, Brno, Czech Republic*, pages 1–9, 1998.

[25] M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing (STOC)*, pages 510–519, 2001.

[26] D. Moreno and J. Wooders. Coalition-proof equilibrium. *Games and Economic Behavior*, pages 80–112, 1996.

[27] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Economic Theory*, pages 511–533, 2001.

[28] J. F. Nash. Equilibrium points in n-person games. In *Proceedings of the National Academy of Sciences*, volume 36, pages 48–49, 1950.

[29] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of ACM*, 49(2):236–259, 2002.

[30] H. Sachs. Regular graphs with given girth and restricted circuits. *J. London Math. Soc.*, 38:423–429, 1963.

[31] M. Shalom and S. Zaks. A $10/7 + \epsilon$ approximation scheme for minimizing the number of adms in sonet rings. In *First Annual International Conference on Broadband Networks, San-José, California, USA*, October 2004.

[32] L.S. Shapley. The value of n-person games. *Contributions to the theory of games*, pages 31–40, Princeton University Press, 1953.