Stackelberg Strategies for Network Design Games

Angelo Fanelli, Michele Flammini, and Luca Moscardelli

Abstract. We consider the network-design game introduced by Anshelevich et al. in which n source–destination pairs must be connected by n respective players equally sharing the cost of the used links. It is well known that the price of anarchy for this class of games may be as large as n. One approach for reducing this bound is that of resorting to the Stackelberg model, in which for a subset of at most $\lfloor \alpha n \rfloor$ coordinated players, with $0 \le \alpha \le 1$, communication paths inducing better equilibria are fixed. In this paper we show the effectiveness of Stackelberg strategies by providing optimal and nearly optimal bounds on the performance achievable by such Stackelberg strategies. In particular, in contrast to previous works, we are also able to provide Stackelberg strategies computable in polynomial time and lowering the price of anarchy from n to $2\left(\frac{1}{\alpha}+1\right)$. Most of the results are extended to the case in which each player aims at connecting $k > 2$ nodes of the network.

1. Introduction

Congestion games [Rosenthal 73] are a well-established approach to modeling resource sharing among selfish players. In such games, a set of resources is available to a set of n players. Each player possesses a set of strategies, each corresponding to the selection of a subset of the resources. A state of the game is any combination of strategies for the players. The cost incurred by a player in a given state is

defined as the sum of the costs associated with each selected resource, which depends on the number of players choosing that resource. The total cost of a state denotes its quality from a global perspective, which is typically defined as the sum of the players' costs or the maximum among the players' costs. Rosenthal has shown that the natural decentralized mechanism known as Nash dynamics, whereby at each step, some player performs an improvement step switching his strategy to a better alternative, is guaranteed to converge to a pure *Nash equilibrium* [Nash 50], i.e., a fixed point of such dynamics in which no player can perform an improvement step. The Nash equilibrium does not necessarily minimize the total cost. The main tool for quantifying the quality of equilibria and thus the performance degradation due to the players' selfish behavior is the *price of anarchy* (PoA), introduced in [Koutsoupias and Papadimitriou 99], which is formally defined as the worst-case ratio of the total cost of a Nash equilibrium to the optimal total cost.

Network-design games with fair cost allocation, introduced in [Anshelevich et al. 08], are one of the most interesting subclasses of congestion games. In the sequel we will refer to this class as network-design games. In a network-design game we are given an undirected graph with nonnegative costs on the edges and, for each player, a source and a destination node. The goal of each player is to choose a path connecting his source and destination nodes. Thus the edges of the graph correspond to the resources of the game, and the strategy set of each player is given by the set of paths connecting the source and destination nodes associated with the player. The cost of each edge e is shared equally by the set of all players whose selected paths contain e. It is well known that the price of anarchy for this game may be as large as the number of players even for a simple game with two parallel edges.

A few natural approaches for reducing the price of anarchy in noncooperative games have been investigated. An interesting one is the *Stackelberg model* [Korilis et al. 97], which consists in assuming that a central authority exploits a small fraction of coordinated players for improving the quality of the Nash equilibrium reached by the remaining selfish players. The central authority selects a fraction of players, called *coordinated players*, and assigns them to appropriately selected strategies. The algorithm adopted by the authority in selecting the coordinated players and assigning them to strategies is called the *Stackelberg strategy*. Given the strategy for the coordinated players, each of the remaining players, called the *selfish players*, selects his strategy selfishly trying to minimize his cost. The behavior of selfish players leads to a *(Stackelberg) Nash equilibrium* in which none of the selfish players can improve his cost. The goal is to determine an effective Stackelberg strategy that improves the price of anarchy of the game.

Related Work. Network-design with fair cost allocation was introduced in [Anshelevich et al. 08]. In that seminal paper, the authors raised the problem of the bad performance, in terms of price of anarchy, of the game due to the selfish behavior of the players. Motivated by this issue, they started to explore the middle ground between centrally enforced solutions and completely unregulated anarchy by proposing the notion of price of stability, that is, the ratio of the cost of the cheapest Nash equilibrium to the cost of an optimal solution.

The first consideration of the use of the Stackelberg model as a means of improving the performance of a system was in [Korilis et al. 97]. Subsequently, the problem of improving the price of anarchy of noncooperative games by means of Stackelberg strategies was considered in [Roughgarden 04]. Stackelberg strategies have been investigated in the context of congestion games with nondecreasing latency functions. In particular, all previous research focused on the nonatomic setting, for instance, [Bonifaci et al. 08, Kaporis and Spirakis 09, Karakostas and Kolliopoulos 09, Kumar and Marathe 02, Roughgarden 04, Sharma and Williamson 07, Swamy 07], and just recently, atomic congestion games were considered in [Fotakis 07].

With respect to Stackelberg strategies for network-design games, a different Stackelberg model was proposed in [Elias et al. 10] in which the central authority (being the leader) directly shares the cost of some edge of the network with the selfish user, and the authors provided numerical results showing the effectiveness of such an approach. To the best of our knowledge, no work has analytically investigated the effectiveness of Stackelberg strategies for congestion games with decreasing delay functions and in particular for network-design games. Nevertheless, several works have dealt with the problem of improving the price of anarchy of network-design games. In particular, [Chen et al. 08] studied the problem of designing a different mechanism for sharing the cost of each edge to optimize equilibrium behavior, while [Balcan et al. 09] considered a public-service advertising model in which each player follows with constant probability a central authority suggestion (see Remark 3.7 for a more detailed comparison with our results); recently, such a model has been investigated in the context of positioning games [Demaine and Zadimoghaddam 12].

In [Bilò et al. 10], it is shown in a social knowledge framework that the price of anarchy is reduced when the players are unaware of the presence of the other players. Finally, it was observed in [Chekuri et al. 06] that the price of anarchy strongly depends on the initial state from which the players begin play. In particular, it was proved that the price of anarchy strongly decreases when one considers only dynamics starting from "empty" states, that is, states in which no player has selected any strategy. All the results in [Chekuri et al. 06] were subsequently improved in [Charikar et al. 08].

Our Contribution. In this paper, we investigate Stackelberg strategies for network-design games. To the best of our knowledge, this is the first work on Stackelberg strategy for congestion games with decreasing latency functions. In particular, we show the effectiveness of Stackelberg strategies in reducing the price of anarchy by providing optimal and nearly optimal bounds on the performance achievable. More precisely, given a subset of at most $\lfloor \alpha n \rfloor$ coordinated players with $0 < \alpha \le 1$, in the case of a single source node, the price of anarchy becomes $\frac{1}{\alpha} + \frac{1}{2}$. Moreover, in the general multiple-sources case, it is at most $\frac{1}{\alpha} + 1$, which is only a subtle additive constant apart from the lower bound induced by a single source.

In contrast to previous works [Fotakis 07], we finally address the question of the selection of good Stackelberg strategies in polynomial running time. Namely, given polynomial-time ρ-approximation algorithms for the minimization of the social function, we show that it is possible to compute in polynomial time communication paths for the coordinated players inducing prices of anarchy equal to the ones listed in the previous paragraph times ρ. This gives a polynomial-time selection-inducing price of anarchy of at most $1.39 \left(\frac{1}{\alpha} + \frac{1}{2} \right)$ for the single-source case and $2 \left(\frac{1}{\alpha} + 1 \right)$ in the general case, by exploiting the Steiner tree and Steiner forest approximation results in [Byrka et al. 10] and [Goemans and Williamson 95], respectively.

Moreover, faster Stackelberg strategies can be obtained at the cost of a price of anarchy equal to $2 \left(\frac{1}{\alpha} + \frac{1}{2} \right)$ and $\frac{4}{\alpha} + 1$, for the single-source and the general cases, respectively. Finally, many results can be directly generalized to the case in which players correspond to communication requests among $k > 2$ nodes of the network.

The paper is structured as follows. In the next section we define the model and introduce some useful definitions. In Section 3 we show that there exist Stackelberg strategies dramatically lowering the price of anarchy from n to a value proportional to $\frac{1}{\alpha}$. Section 4 is devoted to providing efficient Stackelberg strategies having such properties. Finally, Section 5 provides some extensions of our results to more general models and states some interesting open questions.

2. Model and Definitions

A *network-design game* is defined by

$$\mathcal{G} = \left(N, G = (V, E), (w_e)_{e \in E}, \left((r_i, t_i) \in V^2 \right)_{i \in N} \right),$$

where N is the set of players, G is an undirected graph having for each edge $e \in E$ a positive cost w_e and each player $i \in N$ has a pair of nodes $(r_i, t_i) \in V^2$

that he wants to connect, r_i and t_i being the source and the destination nodes, respectively. Notice that if r_i is the same node for every player, we are in the special case of a *single-source* network-design game. Let Σ_i denote the strategy set of player i, with a strategy $s_i \in \Sigma_i$ of i consisting of a path connecting r_i and t_i. Let $S = (s_1, s_2, \ldots, s_n) \in \times_{i \in N} \Sigma_i$ be the strategy profile (state) in which player i chooses his strategy $s_i \in \Sigma_i$. We denote by $G(S)$ the subgraph of G comprising all edges used by all players in state S, i.e., $G(S) = \bigcup_{i \in N} s_i$. Given a strategy profile $S = (s_1, \ldots, s_n)$ and an edge $e \in E$, let $n_e(S)$ be the number of players using e in S, i.e., $n_e(S) = |\{i \in N \mid e \in s_i\}|$. We assume that all players using an edge share its cost equally, i.e., for each edge e and each player i using e in state S, the cost charged to player i for e is $c_i^e(S) = \frac{w_e}{n_e(S)}$. The total cost incurred by player i in S is defined as the sum of the shared costs of all edges used by i, i.e.,

$$c_i(S) = \sum_{e \in s_i} c_i^e(S) = \sum_{e \in s_i} \frac{w_e}{n_e(S)}.$$

The social cost of a strategy profile S is defined as the sum of all the players' costs, i.e., $C(S) = \sum_{i \in N} c_i(S)$. Obviously, $C(S) = \sum_{e \in G(S)} w_e$, that is, the cost of all the edges used by the players in S. An optimal strategy profile, which we denote by OPT, is one with minimum social cost.

Obviously, if S^* is a strategy profile minimizing the social function, then $G(S^*)$ denotes an optimal Steiner forest, that is, a forest with minimum cost connecting all nodes $\{r_i, t_i\}_{i \in N}$.

An α-*Stackelberg algorithm* or α-*Stackelberg strategy*[1] is an algorithm performed by a centralized authority that selects a subset $M \subseteq N$ of $m \leq \lfloor \alpha n \rfloor$ players, with $\alpha \in (0, 1)$, called *coordinated players*, and assigns them to determined strategies. We denote by $\text{Stack}(i)$ the strategy assigned to player $i \in M$. Thus, given an instance of the game, the output of a Stackelberg strategy is a pair $(M \subseteq N, (\text{Stack}(i))_{i \in M})$.

In order to obtain results for α that are as general as possible, i.e., having dependence on n, α can also be considered a function of n, i.e., $\alpha : \mathbb{N} \to (0, 1)$; when it is worth underlining that α is a function of n, we will denote it by $\alpha(n)$. The subset of players $N \setminus M$ is the set of *selfish players*. Each player in $N \setminus M$ acts selfishly and attempts to choose the strategy lowering his cost,

---

[1] Note that the term *strategy* here does not refer to the strategies available to players. We adopt such a term for indicating a Stackelberg algorithm because it is widely used in the literature; however, in order to avoid confusion, throughout the paper we will always call a strategy that is a Stackelberg algorithm by the name *Stackelberg strategy*.

given the strategy choices of other players. Given a strategy profile S and a strategy $s_i' \in \Sigma_i$, let $(S \oplus s_i') = (s_1, \ldots, s_{i-1}, s_i', s_{i+1}, \ldots, s_n)$, i.e., the strategy profile obtained from S if player i changes his strategy from s_i to s_i'. A state S is a *Stackelberg Nash equilibrium* if for every player $i \in N \setminus M$ and strategy $s_i' \in \Sigma_i$, we have $c_i(S) \leq c_i(S \oplus s_i')$, i.e., no player in $N \setminus M$ can improve his individual cost by unilaterally changing his strategy. The *price of anarchy* (PoA) is the ratio $C(S)/\text{OPT}$, where S is the Stackelberg Nash equilibrium of maximum social cost. The goal is to design Stackelberg strategies able to lower the price of anarchy.

3. Existence of "Good" Stackelberg Strategies

In this section, we prove the existence of "good" Stackelberg strategies, i.e., Stackelberg strategies lowering the price of anarchy from n to a value proportional to $\frac{1}{\alpha}$, for network-design games. We first prove a lower bound on the performance of every α-Stackelberg strategy.

Theorem 3.1. *For fixed $0 < \delta < 1$, there exist $\alpha < \delta$ and a (single-source) network-design game for which no α-Stackelberg strategy inducing a price of anarchy lower than $\frac{1}{\alpha} + \frac{1}{2}$ exists.*

Proof. Let ϵ be a positive real number. Let $n = \lceil \frac{2}{\delta} \rceil$, $\delta' = \frac{2}{n} \leq \delta$, and $\alpha = \delta' - \epsilon$. Consider the instance depicted in Figure 1, where all the n players try to connect node t to r. They can choose among three different edges: the top one with cost 1, the bottom one with cost $x = \frac{1}{\delta'} - \frac{1}{2} \geq 1$, and the edge with cost n.

The optimal strategy profile is the one in which every player chooses the edge with cost 1, and $\text{OPT} = 1$.

The α-Stackelberg strategy can coordinate at most $\lfloor \alpha \frac{2}{\delta'} \rfloor = \lfloor \frac{2\alpha}{\alpha + \epsilon} \rfloor = 1$ players; clearly, if it does not coordinate any player, than a Nash equilibrium of cost $n \geq 1 + x$ exists in which all the players select the edge of cost n. If the Stackelberg
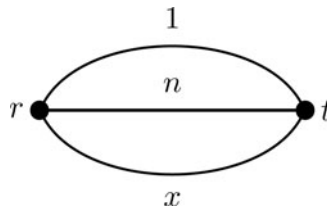
Figure 1. Lower bound on the price of anarchy for a Stackelberg strategy.

strategy fixes the coordinated player on the top edge, there exists an equilibrium in which all the other $n-1$ players select the bottom edge. In fact, in such a way, every selfish player pays $\frac{x}{n-1} = \frac{1}{2}$, which is exactly the cost they would experience by switching to the top edge. Therefore, the social cost of the worst equilibrium is $1 + x$.

If the Stackelberg strategy fixes the coordinated player on the bottom edge, it is easy to verify by a similar argument that there exists a Nash equilibrium in which all the other $n-1$ players select the top edge. Therefore, the social cost of the worst equilibrium is again $1 + x$.

Since

$$1 + x = \frac{1}{\delta'} + \frac{1}{2} = \frac{1}{2} + \frac{1}{\alpha + \epsilon},$$

the claim follows by letting ϵ go to 0. $\qquad\square$

Notice that in all the above instances, α has a fixed dependence on n, being roughly equal to $\frac{2}{n}$. The following theorem shows that whatever the real dependence of α on n, a slightly worse lower bound, equal to $\frac{1}{\alpha}$, holds.

Theorem 3.2. *For every function $\alpha : \mathbb{N} \to \left(0, \frac{1}{2}\right)$ such that $\alpha(k) \geq \frac{1}{k}$ $\forall k \in \mathbb{N}$ and n large enough, there exists a (single-source) network-design game for which no $\alpha(n)$-Stackelberg strategy inducing a price of anarchy lower than $\frac{1}{\alpha(n)}$ exists.*

Proof. The construction is very similar to the one exploited in the proof of Theorem 3.1. Consider the instance depicted in Figure 1, where all the n players try to connect node t to r. They can choose between three different edges: the top one with cost 1, the bottom one with cost

$$x = \frac{n - \lfloor n\alpha(n) \rfloor}{\lfloor n\alpha(n) \rfloor + 1},$$

and the edge with cost n; since $\alpha(n) < \frac{1}{2}$, we obtain that $x \geq 1$.

An optimal strategy profile is the one in which every player chooses the edge of cost 1, and $\mathrm{OPT} = 1$.

The $\alpha(n)$-Stackelberg strategy can coordinate at most $\lfloor n\alpha(n) \rfloor$ players; clearly, if it does not coordinate any player, than a Nash equilibrium of cost $n \geq 1 + x$ exists in which all the players select the edge of cost n.

If the Stackelberg strategy fixes all the $m \leq \lfloor n\alpha(n) \rfloor$ coordinated players on the top edge, there exists an equilibrium in which all the other $n-m$ players select the bottom edge. In fact, in such a way, every selfish player pays $\frac{x}{n-m}$, which is a lower bound to the cost $\frac{1}{m+1}$ they would experience by switching to the top edge. Therefore, the social cost of the worst equilibrium is $1 + x$.

If the Stackelberg strategy fixes the m coordinated players to the bottom edge of cost $x \geq 1$, it is easy to verify by a similar argument that there exists a Nash equilibrium in which all the other $n - m$ players select the top edge. Therefore, the social cost of the worst equilibrium is again $1 + x$. Finally, if the Stackelberg strategy fixes some coordinated players to the top edge and some other coordinated players to the bottom edge, it trivially holds that the social cost of any equilibrium is $1 + x$.

Since

$$1 + x = \frac{n + 1}{\lfloor n\alpha(n) \rfloor + 1},$$

the claim follows for n large enough. □

A natural class of α-Stackelberg strategies that we define is the (α, β, S)-*deterministic scale* (DS) class with $\alpha, \beta \in (0, 1)$ and $S \in \times_{i \in N} \Sigma_i$, which, given a configuration S, selects a set $M \subseteq N$ of $m \leq \lfloor \alpha n \rfloor$ players and assigns to them the same strategies they use in S in such a way that for every resource $e \in E$, at least $\lfloor \beta n_e(S) \rfloor$ players in M use e.

Lemma 3.3. *The* PoA *induced by any α-Stackelberg strategy in the (α, β, S)-DS class is at most $\rho\left(\frac{1}{\beta} + \frac{1}{2}\right)$, where ρ is equal to $\frac{C(S)}{\text{OPT}}$.*

Proof. Let m_e be the number of players fixed by the considered Stackelberg strategy on edge e. We partition the set of edges $G(S)$ used in the solution $S = (s_1, \ldots, s_n)$ into three sets, $E_{\text{light},0}$, $E_{\text{light},\geq 1}$ and E_{heavy} such that:

- $E_{\text{light},0}$ contains all and only the edges e such that $\lfloor \beta n_e(S) \rfloor = 0$ and $m_e = 0$, i.e., the edges not used by the coordinated players;

- $E_{\text{light},\geq 1}$ contains all and only the edges e such that $\lfloor \beta n_e(S) \rfloor = 0$ and $m_e \geq 1$;

- E_{heavy} contains all and only the edges e such that $\lfloor \beta n_e(S) \rfloor \geq 1$ (and therefore also $m_e \geq \lfloor \beta n_e(S) \rfloor \geq 1$).

Moreover, let

$$C_{\text{light},0}(S) = \sum_{e \in E_{\text{light},0}} w_e, \quad C_{\text{light},\geq 1}(S) = \sum_{e \in E_{\text{light},\geq 1}} w_e,$$

$$C_{\text{heavy}}(S) = \sum_{e \in E_{\text{heavy}}} w_e,$$

that is, the cost of the solution S with respect to edges in $E_{\text{light},0}$, $E_{\text{light},\geq 1}$, and E_{heavy}, respectively. Since $E_{\text{light},0}$, $E_{\text{light},\geq 1}$, and E_{heavy} constitute a partition

of $G(S)$, we have $C_{\text{light},0}(S) + C_{\text{light},\geq 1}(S) + C_{\text{heavy}}(S) = C(S)$. Finally, let $\hat{S}$ be any equilibrium.

We now focus on the selfish players. For every player $i \in N \setminus M$,

$$
\begin{aligned}
c_i(\hat{S}) &\leq \sum_{e \in s_i} \frac{w_e}{m_e + 1} = \sum_{e \in s_i} \frac{c_i^e(S) n_e(S)}{m_e + 1} \\
&\leq \sum_{e \in s_i \cap E_{\text{heavy}}} \frac{c_i^e(S) n_e(S)}{m_e + 1} + \sum_{e \in s_i \cap E_{\text{light},0}} c_i^e(S) n_e(S) + \sum_{e \in s_i \cap E_{\text{light},\geq 1}} \frac{c_i^e(S) n_e(S)}{2} \\
&\leq \sum_{e \in s_i \cap E_{\text{heavy}}} \frac{c_i^e(S) n_e(S)}{\beta n_e(S)} + \sum_{e \in s_i \cap E_{\text{light},0}} \frac{c_i^e(S) n_e(S)}{\beta n_e(S)} + \sum_{e \in s_i \cap E_{\text{light},\geq 1}} \frac{c_i^e(S) n_e(S)}{2\beta n_e(S)} \\
&= \frac{1}{\beta} \sum_{e \in s_i \cap (E_{\text{light},0} \cup E_{\text{heavy}})} c_i^e(S) + \frac{1}{2\beta} \sum_{e \in s_i \cap E_{\text{light},\geq 1}} c_i^e(S).
\end{aligned}
$$

Since by the definition of edges in E_{heavy}, at every edge $e \in E_{\text{heavy}}$ the ratio between the number of selfish players and the total number of players in solution S is at most $\frac{n_e(S) - \lfloor \beta n_e(S) \rfloor}{n_e(S)} \leq 1 - \frac{\beta}{2}$, we obtain that

$$
\begin{aligned}
\sum_{i \in N \setminus M} c_i(\hat{S}) &\leq \frac{1}{\beta} \sum_{i \in N \setminus M} \sum_{e \in s_i \cap E_{\text{heavy}}} c_i^e(S) + \frac{1}{\beta} \sum_{i \in N \setminus M} \sum_{e \in s_i \cap E_{\text{light},0}} c_i^e(S) \\
&\quad + \frac{1}{2\beta} \sum_{i \in N \setminus M} \sum_{e \in s_i \cap E_{\text{light},\geq 1}} c_i^e(S) \\
&\leq \left(\frac{1}{\beta} - \frac{1}{2} \right) C_{\text{heavy}}(S) + \frac{1}{\beta} C_{\text{light},0}(S) + \frac{1}{2\beta} C_{\text{light},\geq 1}(S).
\end{aligned}
$$

On the other hand, since only edges in E_{heavy} and in $E_{\text{light},\geq 1}$ can be used by the controlled players, we have that

$$
\sum_{i \in M} c_i(\hat{S}) \leq C_{\text{heavy}}(S) + C_{\text{light},\geq 1}(S).
$$

By summing over all players, we finally obtain

$$
\begin{aligned}
C(\hat{S}) &= \sum_{i \in N \setminus M} c_i(\hat{S}) + \sum_{i \in M} c_i(\hat{S}) \\
&\leq \left(\frac{1}{\beta} + \frac{1}{2} \right) C_{\text{heavy}}(S) + \frac{1}{\beta} C_{\text{light},0}(S) + \left(\frac{1}{2\beta} + 1 \right) C_{\text{light},\geq 1}(S) \\
&\leq \left(\frac{1}{\beta} + \frac{1}{2} \right) C(S),
\end{aligned}
$$

because for every $0 \leq \beta \leq 1$, we have $\frac{1}{2\beta} + 1 \leq \frac{1}{\beta} + \frac{1}{2}$. $\square$

Algorithm 1. DS for single-source network-design games.

1: **procedure** DS-SS(state $S = (s_1, s_2, \ldots, s_n)$ such that $G(S)$ is a tree)
2:      $M \leftarrow \varnothing$
3:      Let $T = G(S)$ be the tree induced by S
4:      Visit the edges of T in reverse order with respect to breadth-first search
5:      **for** every visited edge e **do**
6:          Let $n'_e(S) = |\{i \in M \mid e \in s_i\}|$
7:          Let Q be a subset of $\lfloor \alpha n_e(S) \rfloor - n'_e(S)$ players in $N \setminus M$ using edge e
8:          $M \leftarrow M \cup Q$
9:      **end for**
10:      For all $i \in M$, $\text{Stack}(i) = s_i$
11: **end procedure**

We first show that there exists a Stackelberg strategy in the DS class for the single-source network-design game.

Theorem 3.4. *For the single-source network-design game, the α-Stackelberg strategy defined by Algorithm 1 on input S, with $G(S)$ a tree, belongs to the (α, α, S)-DS class.*

Proof. In order to show that the α-Stackelberg strategy defined by Algorithm 1 on input S belongs to the (α, α, S)-DS class, we have to verify two properties:

1. $\mid M \mid \leq \lfloor \alpha n \rfloor$.

2. For every resource $e \in E$, at least $\lfloor \alpha n_e(S) \rfloor$ players in M use e.

Property 2 is trivially satisfied by the algorithm. In fact, whenever an edge e is processed, the algorithm adds an additional $\lfloor \alpha n_e(S) \rfloor - n'_e(S)$ players to the set of $n'_e(S)$ players using e already in M (lines 6–8).

In order to verify property 1, since all the paths have the source as an endpoint, we have that for each edge e, the number $n_e(S)$ of paths crossing it is equal to the number of all paths that cross at least an edge of the subtree rooted at the endpoint of e farthest from the root. Therefore, if we consider all the edges $e_1, \ldots, e_k$ having the root as an endpoint, the sum of the number of paths crossing them is n, i.e., $\sum_{j=1}^{k} n_{e_j}(S) = n$. Since the coordinated players on edge e_j $(j = 1, \ldots, k)$ number exactly $\lfloor \alpha n_{e_j}(S) \rfloor$, we obtain that

$$m = \sum_{j=1}^{k} \lfloor \alpha n_{e_j}(S) \rfloor \leq \left\lfloor \alpha \sum_{j=1}^{k} n_{e_j}(S) \right\rfloor = \lfloor \alpha n \rfloor,$$

and the proof is complete. $\square$
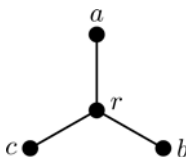
Figure 2. The subgraph of G induced by an optimal solution.

Therefore, the following corollary follows as an immediate consequence of Theorem 3.1 combined with Lemma 3.3 and Theorem 3.4.

Corollary 3.5. *In the single-source network-design game, let S^* denote a strategy profile such that $C(S^*) = \mathrm{OPT}$. The α-Stackelberg strategy defined by Algorithm 1 on input S^* is optimal, i.e., it induces games with* PoA *at most $\frac{1}{\alpha} + \frac{1}{2}$.*

Unfortunately, a Stackelberg strategy belonging to the (α, α, S)-DS class cannot be obtained for general (not only single-source) network-design games. In fact, by means of Figure 2, we show that it is not always possible to control at least $\lfloor \alpha n_e(S) \rfloor$ players for each resource e by controlling at most $\lfloor \alpha n \rfloor$ players in the whole instance. In Figure 2 we consider the subgraph $G(S)$ induced by a state S in a game with three players having requests (a, b), (b, c), and (c, a), respectively. Every edge is used by two players. Thus if we consider $\alpha = \frac{1}{2}$, we must select a subset of the players such that every edge is covered by at least one of them. In order to do that, it is easy to see that we must pick at least two out of the three players, that is, more than half of the total number of players.

In order to obtain an optimal deterministic Stackelberg strategy for the general case, let us introduce another α-Stackelberg strategy, the (α, S)-*probabilistic scale* (PS) Stackelberg strategy, which controls $\lfloor \alpha n \rfloor$ players selected uniformly randomly and assigns them the same strategy they use in solution S. In the following, random variables are denoted by boldface letters.

Theorem 3.6. *In the network-design game, the (α, S)-PS Stackelberg strategy induces games with expected* PoA *at most $\rho \left(\frac{1}{\alpha} + 1 \right)$, where $\rho = \frac{C(S)}{\mathrm{OPT}}$.*

Proof. Let $\hat{S}$ be any Nash equilibrium for the game. We are interested in bounding from above the expected cost in $\hat{S}$ of every player $i \in N \setminus M$. In fact, players in M pay at most the cost of the initial solution S, i.e., $\sum_{i \in M} c_i(\hat{S}) \leq C(S)$.

We will consider the following random variables, depending on the selected subset of $\lfloor \alpha n \rfloor$ coordinated players, corresponding to already defined quantities:

- $\mathbf{c}_i(\hat{S})$, measuring the cost paid by player i in solution $\hat{S}$;

- $\mathbf{c}_i^e(\hat{S})$, measuring the cost paid by player i over edge e, in solution $\hat{S}$;

- $\mathbf{C}(\hat{S})$, measuring the social cost of solution $\hat{S}$.

We first bound $\mathbb{E}\left[\mathbf{c}_i^e(\hat{S})\right]$ from above for every $i \in N \setminus M$. To this end, we introduce the random variable $\mathbf{a}_i^e(\hat{S})$ indicating how much player i would pay on edge e if only coordinated players (and himself) used such an edge. Clearly, $\mathbb{E}\left[\mathbf{c}_i^e(\hat{S})\right] \le \mathbb{E}\left[\mathbf{a}_i^e(\hat{S})\right]$. Since $\Pr\left(\mathbf{a}_i^e(\hat{S}) = \frac{w_e}{x+1}\right)$ induces a hypergeometric probability distribution, we have

$$\Pr\left(\mathbf{a}_i^e(\hat{S}) = \frac{w_e}{x+1}\right) = \frac{\binom{n_e(S)}{x}\binom{n-n_e(S)}{\lfloor \alpha n \rfloor - x}}{\binom{n}{\lfloor \alpha n \rfloor}},$$

and we obtain

$$
\begin{aligned}
\mathbb{E}\left[\mathbf{c}_i^e(\hat{S})\right] &\le \mathbb{E}\left[\mathbf{a}_i^e(\hat{S})\right] \\
&\le w_e \sum_{x=0}^{\min\{\lfloor \alpha n \rfloor, n_e(S)\}} \frac{\Pr\left(\mathbf{a}_i^e(\hat{S}) = \frac{w_e}{x+1}\right)}{x+1} \\
&= \frac{w_e}{\binom{n}{\lfloor \alpha n \rfloor}} \sum_{x=0}^{\lfloor \alpha n \rfloor} \frac{\binom{n_e(S)}{x}\binom{n-n_e(S)}{\lfloor \alpha n \rfloor - x}}{x+1} \qquad (3.1) \\
&= \frac{w_e}{\binom{n}{\lfloor \alpha n \rfloor}(n_e(S)+1)} \sum_{x=0}^{\lfloor \alpha n \rfloor} \binom{n_e(S)+1}{x+1}\binom{n-n_e(S)}{\lfloor \alpha n \rfloor - x} \\
&= \frac{w_e}{\binom{n}{\lfloor \alpha n \rfloor}(n_e(S)+1)} \sum_{y=1}^{\lfloor \alpha n \rfloor + 1} \binom{n_e(S)+1}{y}\binom{n-n_e(S)}{\lfloor \alpha n \rfloor + 1 - y} \\
&\le \frac{w_e \binom{n+1}{\lfloor \alpha n \rfloor + 1}}{\binom{n}{\lfloor \alpha n \rfloor}(n_e(S)+1)} \qquad (3.2) \\
&= \frac{w_e(n+1)}{(\lfloor \alpha n \rfloor + 1)(n_e(S)+1)} \\
&\le \frac{w_e n}{\alpha n n_e(S)} = \frac{w_e}{\alpha n_e(S)},
\end{aligned}
$$

where equality (3.1) holds because if $\min\{\lfloor \alpha n \rfloor, n_e(S)\} = n_e(S)$, then for all $x = n_e(S) + 1, \ldots, \lfloor \alpha n \rfloor$, $\binom{n_e(S)}{x} = 0$ by the definition of binomial coefficients, and inequality (3.2) holds by Vandermonde's identity.

By summing over all players $i \in N$, since $\sum_{i \in M} c_i(\hat{S}) \le C(S)$, we obtain

$$\mathbb{E}\left[\mathbf{C}(\hat{S})\right] = \sum_{i \in M} \mathbb{E}\left[\mathbf{c}_i(\hat{S})\right] + \sum_{i \in N \setminus M} \mathbb{E}\left[\mathbf{c}_i(\hat{S})\right] \le C(S) + \sum_{i \in N \setminus M} \sum_{e \in s_i} \mathbb{E}\left[\mathbf{c}_i^e(\hat{S})\right]$$

$$\le C(S) + \sum_{i \in N \setminus M} \sum_{e \in s_i} \frac{w_e}{\alpha n_e(S)} = C(S) + \sum_{i \in N \setminus M} \sum_{e \in s_i} \frac{c_i^e(S) n_e(S)}{\alpha n_e(S)}$$

$$= C(S) + \frac{1}{\alpha} \sum_{i \in N \setminus M} \sum_{e \in s_i} c_i^e(S) = C(S) + \frac{1}{\alpha} \sum_{i \in N \setminus M} c_i(S)$$

$$\le C(S) + \frac{1}{\alpha} C(S) = \left(1 + \frac{1}{\alpha}\right) C(S).$$

The claim follows by recalling that $C(S) \le \rho \cdot \text{OPT}$. $\square$

Remark 3.7. The result of Theorem 3.6 nicely relates to a result in [Balcan et al. 09], in which the authors study a "public-service advertising" model in which a centralized authority suggests to each player a strategy, and each player follows such a suggestion with a given probability α. The main difference is that in [Balcan et al. 09], the authority controls an expected number of $\alpha \cdot n$ players, while in our model we always control a set of $\lfloor \alpha n \rfloor$ players, and such a set is selected uniformly at random among all the possible sets of $\lfloor \alpha n \rfloor$ players. Interestingly, both analyses lead to similar results concerning the expected value of the random variable $\mathbf{a}_i^e$ introduced in the proof of Theorem 3.6, even if the former model deals with a binomially distributed variable, while in the latter model we deal with a hypergeometric probability distribution. Finally, it is worth noticing that our result on the expected PoA, being at most $\rho\left(\frac{1}{\alpha} + 1\right)$, is tighter than the corresponding one in [Balcan et al. 09], being $O\left(\frac{\rho}{\alpha}\right)$; in fact, the authors of [Balcan et al. 09] provide a proof in which it is shown that the expected PoA is at most $c\frac{\rho}{\alpha}$, with the constant c approaching 3 as the number of players grows. Such a better result has been obtained by a cleverer analysis, which takes into account the cost incurred by the controlled players, i.e., by accounting at most ρ for the cost paid by the controlled players.

Corollary 3.8. *In the network-design game, let S^* denote a strategy profile such that $C(S^*) = \text{OPT}$; the (α, S^*)-PS Stackelberg strategy induces games with expected PoA at most $\frac{1}{\alpha} + 1$.*

Moreover, since the expected PoA is at most $\frac{1}{\alpha} + 1$, there must exist a deterministic Stackelberg strategy with PoA at most the expected one, and the following corollary holds.

Corollary 3.9. *In the network-design game there exists an almost optimal deterministic α-Stackelberg strategy inducing games with* PoA *at most $\frac{1}{\alpha} + 1$.*

Let us note some interesting points. First of all, let us introduce a useful definition: Consider an α-Stackelberg strategy and let $\hat{S}$ be any Nash equilibrium of the network-design game induced by it. We say that such an α-Stackelberg strategy is *fair* with respect to a state S if for every player $i \in N$, we have $c_i(\hat{S}) = O\left(\frac{c_i(S)}{\alpha}\right)$. Notice that if an α-Stackelberg strategy is *fair* with respect to state S, it is able to induce Nash equilibria with price of anarchy $O\left(\frac{\rho}{\alpha}\right)$, where $\rho = \frac{C(S)}{\text{OPT}}$, not only for the sum social function considered in this paper, but also with respect to a wider class of social functions, including, for instance, the one considering the maximum cost among all the player costs (see Section 5 for further details).

Even if Corollary 3.9 shows the existence of an almost optimal deterministic α-Stackelberg strategy in a nonconstructive way, in Section 4, we will show that by exploiting derandomization techniques, it is possible to obtain a polynomial-time computable deterministic α-Stackelberg strategy inducing the claimed price of anarchy.

Clearly, such an almost optimal deterministic α-Stackelberg strategy also holds for the particular case of a single-source network-design-game; nevertheless, we have presented the α-Stackelberg strategy of Algorithm 1 not only because of the better additive constant exactly matching the lower bound of Theorem 3.1; in fact, as will be discussed in Section 4, the computational cost of Algorithm 1 is much lower than that of the α-Stackelberg strategy described in the proof of Theorem 4.3, being the α-Stackelberg strategy whose existence is claimed in Corollary 3.9. Moreover, as stated in the following theorem, α-Stackelberg strategies in the (α, β, S)-DS class ensure a stronger property, i.e., they are fair with respect to S.

Taking in mind such observations, as a further contribution, in Section 4, at the cost of a multiplicative factor equal to 2 in the price of anarchy, a fair and faster α-Stackelberg strategy for the general network-design game will be presented.

Theorem 3.10. *Every α-Stackelberg strategy belonging to the (α, β, S)-DS class, with $\beta = \Omega(\alpha)$, is* fair *with respect to S.*

Proof. Let $\hat{S}$ be any Nash equilibrium of the network-design game induced by the considered α-Stackelberg strategy. As shown in the proof of Lemma 3.3, for every player $i \in N \setminus M$,

$$c_i(\hat{S}) \leq \frac{1}{\beta} \sum_{e \in s_i \cap (E_{\text{light},0} \cup E_{\text{heavy}})} c_i^e(S) + \frac{1}{2\beta} \sum_{e \in s_i \cap E_{\text{light},\geq 1}} c_i^e(S) \leq \frac{c_i(S)}{\beta}$$

$$= O\left(\frac{c_i(S)}{\alpha}\right).$$

It remains to show that for every coordinated player $i \in M$, $c_i(\hat{S}) = O\left(\frac{c_i(S)}{\alpha}\right)$. Since $\lfloor \beta n_e(S) \rfloor \geq \frac{\beta}{2} n_e(S)$ for every edge used by at least one coordinated player, we obtain

$$c_i(\hat{S}) \leq \sum_{e \in s_i} \frac{w_e}{n_e(\hat{S})} \leq \sum_{e \in s_i} \frac{c_i^e(S) n_e(S)}{\lfloor \beta n_e(S) \rfloor} \leq \sum_{e \in s_i} \frac{c_i^e(S) n_e(S)}{\frac{\beta}{2} n_e(S)} = \frac{2}{\beta} \sum_{e \in s_i} c_i^e(S)$$

$$= \frac{2}{\beta} c_i(S) = O\left(\frac{c_i(S)}{\alpha}\right),$$

which completes the proof. $\square$

4. Efficient Stackelberg Strategies

In this section we focus on the efficient computation of good Stackelberg strategies for network-design games.

Most of the results of Section 3 can be exploited to obtain efficient Stackelberg strategies that although suboptimal, are able to lower the price of anarchy to a value proportional to $\frac{1}{\alpha}$.

The basic idea is that of considering a ρ-approximation $\bar{S}$ instead of an optimal solution S^*. Such an approximate solution is known to be efficiently computable, since the social optimization problem is basically a minimum Steiner forest problem, and a constant approximation is easily obtainable.

In fact, given an edge-weighted graph $G = (V, E)$ and a set of nodes $\{s_i, t_i\}$, $s_i, t_i \in V$, the Steiner forest problem consists in finding a minimum-cost forest such that both nodes of the instance are connected. In [Goemans and Williamson 95] a 2-approximation algorithm for such a problem has been provided, running in time $O(|V|^2 \ln |V|)$. Let $\bar{S}$ denote such a 2-approximate solution.

By Lemma 3.3 and Theorem 3.4, since, as is easily verified, Algorithm 1 has a worst-case time complexity $O(n|E|)$, the following proposition holds.

Proposition 4.1. *In the single-source network-design game, the α-Stackelberg strategy defined by Algorithm 1 on input $\bar{S}$ induces games with* PoA *at most $\frac{2}{\alpha} + 1$. The time complexity of computing such an α-Stackelberg strategy is $O(n|E| + |V|^2 \ln |V|)$.*

Notice that by exploiting the 1.39-approximation algorithm for the Steiner tree problem of [Byrka et al. 10], the PoA can be lowered to $1.39\left(\frac{1}{\alpha} + \frac{1}{2}\right)$, but the running time of such an approximation algorithm increases rapidly as the guaranteed ratio approaches 1.39.

Now we turn our attention to the (α, S)-PS probabilistic Stackelberg strategy. The following proposition is an immediate consequence of Theorem 3.6.

Proposition 4.2. *In the network-design game, the $(\alpha, \bar{S})$-PS is a probabilistic polynomial Stackelberg strategy controlling $\lfloor \alpha n \rfloor$ players and inducing games with expected* PoA *at most $\frac{2}{\alpha} + 2$. The time complexity of computing such a probabilistic α-Stackelberg strategy is $O(|V|^2 \ln |V|)$.*

Starting from the (α, S)-PS Stackelberg strategy and applying standard *derandomization* arguments (in particular, the method of conditional probabilities, as Algorithm 2 does), it is possible to obtain a new deterministic α-Stackelberg strategy for network-design games realizing the properties claimed in Corollary 3.9.

We perform the derandomization with respect to a random variable $\mathbf{C}$ that is an upper bound to the cost of any Nash equilibrium. In particular, such an upper bound is given by the costs $\mathbf{a}_i^e$, $e \in s_i$, that the player i would experience on the resources he selects in S, assuming that the costs of such resources are shared only with the coordinated players: $\mathbf{C} = \sum_{i \in N} \sum_{e \in s_i} \mathbf{a}_i^e$.

Theorem 4.3. *For network-design games, the (deterministic) α-Stackelberg strategy defined by Algorithm 2 induces games with* PoA *at most $\rho\left(\frac{1}{\alpha} + 1\right)$, where $\rho = \frac{C(S)}{\mathrm{OPT}}$.*

Proof. By the proof of Theorem 3.6, we know that the expected value of $\mathbf{C}$ is at most $\left(\frac{1}{\alpha} + 1\right) C(S)$. Therefore, there must exist a deterministic choice of $d_1, \ldots, d_n$ inducing a worst-case equilibrium $\hat{S}$ such that

$$C(\hat{S}) \leq \mathbf{C} \leq \rho\left(\frac{1}{\alpha} + 1\right) \mathrm{OPT}.$$

It is easy to check that at every iteration of the FOR block, d_i is always chosen in such a way that the expected value of $\mathbf{C}$ is at most $\left(\frac{1}{\alpha} + 1\right) C(S)$. Notice that

---

Algorithm 2. Derandomization.

1: **procedure** DERAND(state $S = (s_1, s_2, \dots, s_n)$)
2:     Let $q(d_1, \dots, d_i)$ be the conditional expectation $\mathbb{E}\left[\mathbf{C} \mid (j \in M \Leftrightarrow d_j = 1) \wedge \text{Stack}(j) = s_j, 1 \le j \le i\right]$
3:     **for** $i \leftarrow 1, n$ **do**
4:         **if** $q(d_1, \dots, d_{i-1}, 0) < q(d_1, \dots, d_{i-1}, 1)$ **then**
5:             $d_i \leftarrow 0$
6:         **else**
7:             $d_i \leftarrow 1$
8:         **end if**
9:     **end for**
10:     Put in M all and only the players i for which $d_i = 1$
11:     For all $i \in M$, $\text{Stack}(i) = s_i$
12: **end procedure**

---

this is always possible because for every $i = 1, \dots, n$, starting from a situation in which the expected value of $\mathbf{C}$ conditioned to the choice of $d_1, \dots, d_{i-1}$ is at most x, there must exist a choice for d_i ($d_i = 0$ or $d_i = 1$) for which the expected value of $\mathbf{C}$, conditioned also by the new choice of d_i, is at most x. $\square$

Notice that the time complexity of Algorithm 2 depends on the computational cost of the conditional expectations q. In order to compute such expectations, techniques similar to the one exploited in the proof of Theorem 3.6 can be used, i.e., letting $\gamma \equiv (i \in M \Leftrightarrow d_i = 1) \wedge \text{Stack}(i) = s_i$, $D = |\{i \in N \mid d_i = 1\}|$, and $D_e = |\{i \in N \mid d_i = 1 \wedge e \in \text{Stack}(i)\}|$,

$$
\mathbb{E}\left[\mathbf{C} \mid \gamma\right] = \sum_{i \in N} \sum_{e \in s_i} \mathbb{E}\left[\mathbf{a}_i^e \mid \gamma\right]
$$

$$
= \sum_{i \in M} \sum_{e \in s_i} \sum_{x=0}^{\lfloor \alpha n \rfloor - D} \frac{w_e}{D_e + x} \Pr\left(\mathbf{a}_i^e = \frac{w_e}{D_e + x} \mid \gamma\right)
$$

$$
+ \sum_{i \in N \setminus M} \sum_{e \in s_i} \sum_{x=0}^{\lfloor \alpha n \rfloor - D} \frac{w_e}{D_e + x + 1} \Pr\left(\mathbf{a}_i^e = \frac{w_e}{D_e + x + 1} \mid \gamma\right),
$$

where for every $i \in M$ and $j \in N \setminus M$,

$$
\Pr\left(\mathbf{a}_i^e = \frac{w_e}{D_e + x} \mid \gamma\right) = \Pr\left(\mathbf{a}_j^e = \frac{w_e}{D_e + x + 1} \mid \gamma\right)
$$

$$
= \frac{\binom{n_e(S) - D_e}{x} \binom{n - D - n_e(S) + D_e}{\lfloor \alpha n \rfloor - D - x}}{\binom{n - D}{\lfloor \alpha n \rfloor - D}}.
$$

It is easy to verify that the computation of $\mathbb{E}\left[\mathbf{a}_i^e \mid \gamma\right]$ can be performed in time $O(n^2)$. Since the number of edges used in S is $O(|E|)$ and Algorithm 2 has n iterations, its worst-case time complexity is $O(n^3|E|)$.

Therefore, by running Algorithm 2 on input $\bar{S}$, we obtain the following result.

Corollary 4.4. *In the network-design game, the (deterministic) α-Stackelberg strategy induced by Algorithm 2 on input $\bar{S}$ induces games with PoA at most $\frac{2}{\alpha} + 2$. The time complexity of computing such an α-Stackelberg strategy is $O(n^3|E| + |V|^2 \ln |V|)$.*

Finally, in the following, we will show how it is possible to reduce the time complexity of a Stackelberg strategy by paying a factor almost equal to 2 in the PoA.

In particular, we focus on the determination of a (suboptimal) faster α-Stackelberg strategy for the general network-design game.

In order to describe the desired α-Stackelberg strategy, we need some additional definitions: Given a strategy profile $S = (s_1, \ldots, s_n)$ such that $G(S)$ is a tree, we define the split instance of the considered game as the same instance in which we have a new *split player set* N'. In particular, consider the tree $G(S)$ rooted at a generic node u; in order to build the player set N', we split each player (r_i, t_i), $i = 1, \ldots, n$, into at most two players (r_i, v_i) and (t_i, v_i), where v_i is the common ancestor in tree $G(S)$ of r_i and t_i (notice that if $v_i \equiv r_i$ or $v_i \equiv t_i$, then player i is not split). Therefore, $|N| \leq |N'| \leq 2|N|$, and the *split player strategy profile S'* corresponding to S is built by associating to each player (r_i, v_i) (respectively (t_i, v_i)) in N' the strategy corresponding to the unique path connecting r_i and v_i (respectively t_i and v_i) in tree $G(S)$. We denote a generic player in N' by x.

A *fractional set* $\tilde{A}$ of players is a set in which each player may be present only for a fractional portion with value p, $0 < p \leq 1$ (if $p = 1$, the player is completely contained in $\tilde{A}$); notice that different portions of the same player can belong to two disjoint fractional sets. Given a fractional set of players $\tilde{A}$, let $|\tilde{A}|$ be the sum of all the values of the fractional portions belonging to $\tilde{A}$. Given a set of players A, let $\text{frac}(A)$ be the corresponding fractional set of players, in which all the players in A are completely contained in $\text{frac}(A)$, i.e., all the fractional portions in $\text{frac}(A)$ have value 1 and $|\text{frac}(A)| = |A|$. If a fraction p of player x belongs to the set $\tilde{A}$, we write $x^p \in \tilde{A}$. Moreover, $\tilde{B} \subseteq \tilde{A}$ if and only if all the fractional portions of players contained in $\tilde{B}$ are also in $\tilde{A}$.

Given a player $x = (r, t)$, and recalling that u is the root of the tree $G(S)$, let $\text{height}(x) = \min\{\text{dist}(u, r), \text{dist}(u, t)\}$, where $\text{dist}(u, v)$ is defined as the number of edges between u and v in $G(S)$. From now on, we assume that the players

---

Algorithm 3. DS for network-design games.

1: **procedure** DS(state $S = (s_1, s_2, \ldots, s_n)$ such that $G(S)$ is a tree)
2:      $M' \leftarrow \varnothing$
3:      $j \leftarrow 0$
4:      Let $T = G(S)$ be the tree induced by S rooted at a generic node u
5:      Let N' and S' be respectively the *split player set* and the *split player strategy profile* with respect to the tree T rooted at u
6:      $\tilde{N}' \leftarrow \mathrm{frac}(N')$
7:      Visit the edges of T in reverse order with respect to breadth-first search
8:      **for** every visited edge e **do**
9:          Let $N_e(S') = \{x \in N' \mid e \in s'_x\}$ and $n_e(S') = |N_e(S')|$
10:         Let $M_e(S') = \{x \in M' \mid e \in s'_x\}$ and $m_e(S') = |M_e(S')|$
11:         Let $\tilde{N}_e(S')$ be the subset of $\tilde{N}'$ containing all the portions of players x such that $e \in s'_x$
12:         **while** $m_e(S') < \left\lfloor \frac{\alpha}{2} n_e(S') \right\rfloor$ **do**
13:             Let $\tilde{Q} = \mathrm{top}_{\frac{2}{\alpha}}(\tilde{N}_e(S'))$
14:             Let $j = \min\{i \mid x_j^p \in \tilde{Q}, p > 0\}$
15:             $\tilde{N}' \leftarrow \tilde{N}' \setminus \tilde{Q}$
16:             $M' \leftarrow M' \cup \{x_j\}$
17:         **end while**
18:      **end for**
19:      Put in M all the players $i \in N$ such that $(r_i, v_i) \in M'$ or $(t_i, v_i) \in M'$ (or both)
20:      For all $i \in M$, Stack$(i) = s_i$
21: **end procedure**

$x_1, \ldots, x_{|N'|} \in N'$ are numbered in nondecreasing order of height, resolving ties arbitrarily, i.e., $\mathrm{height}(x_1) \leq \mathrm{height}(x_2) \leq \cdots \leq \mathrm{height}(x_{|N'|})$.

Finally, given a fractional set $\tilde{A}$, let $\mathrm{top}_z(\tilde{A}) \subseteq \tilde{A}$ be the set of cardinality $\min\{z, |\tilde{A}|\}$ containing the fractions in $\tilde{A}$ corresponding to the players in N' with the lowest indices. More formally, let $i \leq |N'|$ be the maximum integer such that $|\{x_j^p \mid p > 0, j \leq i, x_j^p \in \tilde{A}\}| \leq z$: $\mathrm{top}_z(\tilde{A})$ contains all the fractions in $\tilde{A}$ corresponding to players x_j with $j \leq i$; moreover, if

$$\delta = \min\{z, |\tilde{A}|\} - \left| \left\{ x_j^p \mid p > 0, j \leq i, x_j^p \in \tilde{A} \right\} \right| > 0,$$

then $\mathrm{top}_z(\tilde{A})$ also contains a (sub)portion δ of the portion of player x_{i+1} belonging to $\tilde{A}$.

Algorithm 3 focuses on the split instance of the game and visits the edges of $G(S)$ in reverse order with respect to breadth-first search. For each visited edge e, it associates to every (fractional) set (containing portions of players selecting edge e in S') of cardinality $\frac{2}{\alpha}$ a different controlled player, i.e., the one with the smallest index having some portion in the considered set.

Theorem 4.5. *For the network-design game, the α-Stackelberg strategy defined by Algorithm 3 on input S belongs to the $(\alpha, \frac{\alpha}{2}, S)$-DS class.*

Proof. We first show that in the auxiliary instance with set player N' and strategy profile S', the $\frac{\alpha}{2}$-Stackelberg strategy that selects players in M' and assigns them their strategy in S' belongs to the $(\frac{\alpha}{2}, \frac{\alpha}{2}, S')$-DS class. To this end, we have to verify two properties:

$1'$ $m' = |M'| \leq \lfloor \frac{\alpha |N'|}{2} \rfloor$.

$2'$ For each resource e, at least $\lfloor \frac{\alpha}{2} n_e(S') \rfloor$ players using resource e in S' are in M'.

First of all, notice that in any iteration of the WHILE statement (line 12), since the edges of T are visited in reverse order with respect to breadth-first search and x_j is selected as the player with minimum height in $G(S)$, the following statements hold:

(i) x_j is such that for all the edges e' visited by the algorithm after e in which a portion belonging to Q_j is present, also $e' \in s'_{x_j}$.

(ii) x_j does not belong to M' when j is chosen at line 14. In fact, no player with some portion in $\tilde{Q}$ belongs to M' at line 14, because whenever a player $x_{j'}$ is selected in some previous iteration of the WHILE statement, since $\frac{2}{\alpha} > 1$ and j' is the minimum index among those players having portions in the considered set, no portions of player $x_{j'}$ can still belong to $\tilde{N'}$ in the following iterations of the WHILE statement.

Therefore, it is always guaranteed that in each iteration of the WHILE statement (line 12) relative to edge e, a subset $\tilde{Q}$ such that $|\tilde{Q}| = \frac{2}{\alpha}$ is present (because otherwise, $\lfloor \frac{\alpha}{2} n_e(S') \rfloor$ subsets of that cardinality would be associated to a player already put in M', and $m_e(S') \geq \lfloor \frac{\alpha}{2} n_e(S') \rfloor$), and also a player x_j can be associated to $\tilde{Q}$.

Property $2'$ is trivially satisfied by the condition of the WHILE statement (line 12).

It remains to show that property $1'$ holds. Since $|N'| = |\tilde{N'}|$ and for every player x_j put in M', a subset $\tilde{Q}$ of $\tilde{N'}$ such that $|\tilde{Q}| = \frac{2}{\alpha}$ is reduced by $\tilde{N'}$, it follows that $|M'| \leq \lfloor \frac{|N'|}{\frac{2}{\alpha}} \rfloor$.

In order to show that the α-Stackelberg strategy defined by Algorithm 3 on input S belongs to the $(\alpha, \frac{\alpha}{2}, S)$-DS class, we have to verify the following:

1. $m \leq \lfloor \alpha n \rfloor$.

2. For each resource e, at least $\lfloor \frac{\alpha}{2} n_e(S) \rfloor$ players using resource e in S are in M.

Property 1 is implied by Property 1′, since $m \leq m'$ and $|N'| \leq 2n$. Property 2 is directly implied by Property 2′, since, as is easily verified, $n_e(S') = n_e(S)$ and for every player $x \in M'$ ($x = l_i$ or $x = r_i$), the corresponding player $i \in N$ is in M.   □

By Lemma 3.3 and Theorem 4.5, since, as is easily verified, Algorithm 3 has a worst-case time complexity $O(n|E|)$, the following proposition holds.

Proposition 4.6. *In the network-design game, the α-Stackelberg strategy defined by Algorithm 3 on input $\bar{S}$ induces games with* PoA *at most $\frac{4}{\alpha} + 1$. The time complexity of computing such an α-Stackelberg strategy is $O(n|E| + |V|^2 \ln |V|)$.*

5. Conclusions and Extensions

In this paper we have given optimal and nearly optimal bounds on the performance achievable by Stackelberg strategies in network-design games.

First of all, it is worth noticing that many results can be directly generalized to the case in which players correspond to communication requests among $k > 2$ nodes of a network. In particular, the probabilistic Stackelberg strategy analyzed in Theorem 3.6 maintains the same performance, and it is easy to modify Algorithm 3 and its analysis in Theorem 4.5 to obtain a deterministic Stackelberg strategy belonging to the $(\alpha, \frac{\alpha}{k}, S)$-DS class. Clearly, all the implications of such results stated in the related corollaries and in Section 4 directly follow.

Furthermore, our probabilistic construction and thus the corresponding bound on the price of anarchy remain valid also in general congestion games with Shapley costs, that is, when resources do not correspond to edges in graphs and strategies to communication paths.

Besides tightening the gaps between the lower and the upper bounds on the price of anarchy in the general multiple-sources case, the following main questions remain open.

The most interesting open question is perhaps that of determining Stackelberg strategies for other social functions such as the MAX social function (i.e., whereby the social cost of a state is given by the maximum cost paid by any player). In this respect, as stated in the preliminary version of this work [Fanelli et al. 10], as a direct consequence of Theorem 3.10, for the MAX social function the PoA induced by any α-Stackelberg strategy in the (α, β, S)-DS class is at most $O(\frac{\rho}{\beta})$, where ρ is the approximation ratio, again with respect to the MAX social function, of state S. Such a result implies that when there exists an optimal solution S^* for the MAX social function such that $G(S^*)$ is a tree, all the results on the existence

of "good" Stackelberg strategies holding for the Stackelberg strategies belonging to the deterministic scale class also hold in the MAX social function case.

The main open problem is that of extending such results to the general case in which such an assumption does not hold, i.e., to the instances of the network-design game in which every optimal state S^* for the MAX social function is such that $G(S^*)$ contains at least a cycle. Furthermore, since for such a social function also the determination of states approximating the social optimum is still an open problem, such results cannot be exploited for determining efficiently computable Stackelberg strategies. Moreover, the analysis of the probabilistic strategy cannot be extended to the MAX social function, because it heavily exploits the *linearity of expectation* property of random variables.

Finally, what about the directed case, in which the graph underlying the network is directed?

References

[Anshelevich et al. 08] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. "The Price of Stability for Network Design with Fair Cost Allocation." *SIAM Journal on Computing* 38:4 (2008), 1602–1623.

[Balcan et al. 09] Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. "Improved Equilibria via Public Service Advertising." In *SODA 2009*, pp. 728–737, 2009.

[Bilò et al. 10] Vittorio Bilò, Angelo Fanelli, Michele Flammini, and Luca Moscardelli. "When Ignorance Helps: Graphical Multicast Cost Sharing Games." *Theoretical Computer Science* 411:3 (2010), 660–671.

[Bonifaci et al. 08] Vincenzo Bonifaci, Tobias Harks, and Guido Schäfer. "Stackelberg Routing in Arbitrary Networks." In *WINE 2008*, pp. 239–250, 2008.

[Byrka et al. 10] Jaroslaw Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. "An Improved LP-Based Approximation for Steiner Tree." In *STOC 2010*, pp. 583–592, 2010.

[Charikar et al. 08] Moses Charikar, Howard J. Karloff, Claire Mathieu, Joseph Naor, and Michael E. Saks. "Online Multicast with Egalitarian Cost Sharing." In *SPAA 2008*, pp. 70–76, 2008.

[Chekuri et al. 06] Chandra Chekuri, Julia Chuzhoy, Liane Lewin-Eytan, Joseph Naor, and Ariel Orda. "Non-cooperative Multicast and Facility Location Games." In *EC 2006*, pp. 72–81, 2006.

[Chen et al. 08] Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. "Designing Networks with Good Equilibria." In *SODA 2008*, pp. 854–863, 2008.

[Demaine and Zadimoghaddam 12] Erik D. Demaine and Morteza Zadimoghaddam. "Constant Price of Anarchy in Network-Creation Games via Public-Service Advertising." *Internet Mathematics* 8 (2012), 29–45.

[Elias et al. 10] Jocelyne Elias, Fabio Martignon, Konstantin Avrachenkov, and Giovanni Neglia. "Socially-Aware Network Design Games." In *INFOCOM 2010*, pp. 41–45, 2010.

[Fanelli et al. 10] Angelo Fanelli, Michele Flammini, and Luca Moscardelli. "Stackelberg Strategies for Network Design Games." In *WINE 2010*, pp. 222–233, 2010.

[Fotakis 07] Dimitris Fotakis. "Stackelberg Strategies for Atomic Congestion Games." In *ESA 2007*, pp. 299–310, 2007.

[Goemans and Williamson 95] Michel X. Goemans and David P. Williamson. "A General Approximation Technique for Constrained Forest Problems." *SIAM Journal on Computing* 24:2 (1995), 296–317.

[Kaporis and Spirakis 09] Alexis C. Kaporis and Paul G. Spirakis. "The Price of Optimum in Stackelberg Games on Arbitrary Single Commodity Networks and Latency Functions." *Theoretical Computer Science* 410 (2009), 745–755.

[Karakostas and Kolliopoulos 09] George Karakostas and Stavros G. Kolliopoulos. "Stackelberg Strategies for Selfish Routing in General Multicommodity Networks." *Algorithmica* 53:1 (2009), 132–153.

[Korilis et al. 97] Yannis A. Korilis, Aurel A. Lazar, and Ariel Orda. "Achieving Network Optima Using Stackelberg Routing Strategies." *IEEE/ACM Transaction on Networking* 5:1 (1997), 161–173.

[Koutsoupias and Papadimitriou 99] Elias Koutsoupias and Christos H. Papadimitriou. "Worst-Case Equilibria." In *STACS 1999*, pp. 404–413, 1999.

[Kumar and Marathe 02] V. S. Anil Kumar and Madhav V. Marathe. "Improved Results for Stackelberg Scheduling Strategies." In *ICALP 2002*, pp. 776–787, 2002.

[Nash 50] John F. Nash. "Equilibrium Points in n-Person Games." *Proceedings of the National Academy of Sciences* 36 (1950), 48–49.

[Rosenthal 73] Robert W. Rosenthal. "A Class of Games Possessing Pure-Strategy Nash Equilibria." *International Journal of Game Theory* 2 (1973), 65–67.

[Roughgarden 04] Tim Roughgarden. "Stackelberg Scheduling Strategies." *SIAM Journal on Computing* 33:2 (2004), 332–350.

[Sharma and Williamson 07] Yogeshwer Sharma and David P. Williamson. "Stackelberg Thresholds in Network Routing Games or the Value of Altruism." In *EC 2007*, pp. 93–102, 2007.

[Swamy 07] Chaitanya Swamy. "The Effectiveness of Stackelberg Strategies and Tolls for Network Congestion Games." In *SODA 2007*, pp. 1133–1142, 2007.

Angelo Fanelli, CNRS (UMR-6211), 19 rue Claude Bloch, 14000 Caen, France (angelo.fanelli@gmail.com)

Michele Flammini, Department of Information Engineering, Computer Science, and Mathematics, University of L'Aquila, Via Vetoio loc. Coppito, 67100 L'Aquila, Italy (michele.flammini@univaq.it)

Luca Moscardelli, Department of Economic Studies, University of Chieti-Pescara, Viale Pindaro 42, 65127 Pescara, Italy (luca.moscardelli@unich.it)