

On Approximate Nash Equilibria in Network Design

Susanne Albers and Pascal Lenzner

Abstract. In this paper, we study a basic network design game in which n self-interested agents, each having individual connectivity requirements, wish to build a network by purchasing links from a given set of edges. A fundamental cost-sharing mechanism is Shapley cost-sharing, which splits the cost of an edge in a fair manner among the agents using the edge. It is well known that in such games, the price of anarchy is n , while the price of stability is $H(n)$, where $H(n)$ denotes the n th harmonic number.

We investigate whether an optimal minimum-cost network represents an attractive, relatively stable state that agents might want to purchase: what extra cost does an agent incur compared to a possible strategy deviation? We employ the concept of α -approximate Nash equilibria, in which no agent can reduce its cost by a factor of more than α . We prove that for single-source games in undirected graphs, every optimal network represents an $H(n)$ -approximate Nash equilibrium. We show that this bound is tight by presenting a matching lower bound. We extend the results to cooperative games in which agents may form coalitions. A combination of strategies forms an α -approximate strong Nash equilibrium if no coalition can deviate such that all members of the coalition reduce their cost by a factor of more than α . We prove that if coalitions of up to c agents are allowed, $1 \leq c \leq n$, then every optimal network represents a $2c(\ln(n/c) + 2)$ -approximate strong Nash equilibrium. We give a corresponding lower bound that is tight up to a small constant factor. Moreover, we extend the results to weighted games and present tight upper and lower bounds on the quality of optimal solutions in noncooperative games. Finally, we show that in general source-sink games and in directed graphs, minimum-cost networks do not represent good states.

1. Introduction

Today, many networks are not built and maintained by a central authority but rather by a large number of economic agents that usually have selfish interests. As a result, game-theoretic approaches for modeling network formation and agent behavior have received considerable research interest over the past years; see, e.g., [Alon et al. 10, Anshelevich et al. 08b, Anshelevich et al. 08a, Chekuri et al. 07, Chen and Roughgarden 10, Chen et al. 10, Corbo and Parkes 05, Demaine et al. 09, Devanur et al. 05, Fabrikant et al. 03, Halevi and Mansour 07, Laoutaris et al. 08, Pal and Tardos 03, Vetta 02].

We study a very basic network design game that has received considerable attention [Albers 09, Anshelevich et al. 08b, Anshelevich et al. 08a, Chekuri et al. 07, Chen and Roughgarden 10, Epstein et al. 09, Fiat et al. 06, Hoefer 09]. Let $G = (V, E, c)$ be a graph, where c is a nonnegative cost function defined on the edges. The graph may be directed or undirected. There are n selfish agents, each having to connect a set of terminals in G . A strategy of an agent i is an edge set connecting the desired terminals. Edges used by the agents have to be paid for. A fundamental cost-sharing mechanism is Shapley cost-sharing, proposed in [Anshelevich et al. 08b] for network design games. In Shapley cost-sharing, the cost of an edge is split in a fair manner among the agents using the edge. More specifically, in an unweighted game, if k agents use an edge e , then each of them pays a share of $c(e)/k$. In a weighted game, each agent i has a positive weight w_i and pays a share proportional to its weight.

Previous work has analyzed stable states in which agents have no incentive to deviate from their strategies. In a standard noncooperative game, a combination of strategies forms a Nash equilibrium if no agent has a better strategy with a strictly smaller cost if all other agents adhere to their strategies. In cooperative games, where coordination among agents is allowed, one is interested in strong Nash equilibria that are resilient to deviations of coalitions of agents [Aumann 59]. A combination of strategies forms a strong Nash equilibrium if there exists no coalition of agents that can jointly change strategy such that every agent of the coalition achieves a strictly smaller cost. There exist two performance measures evaluating Nash equilibria relative to globally optimal solutions. The price of anarchy is the maximum ratio of the total cost incurred by any Nash equilibrium to the cost paid by an optimal solution [Koutsoupias and Papadimitriou 09]. The price of stability is the minimum ratio, i.e., the cost ratio of the best Nash equilibrium relative to the optimum [Anshelevich et al. 08b]. It is shown in [Anshelevich et al. 08b] that for unweighted noncooperative games, the price of anarchy is n , while the price of stability is $H(n)$. Here $H(n) = \sum_{i=1}^n 1/i$

is the n th harmonic number, which is closely approximated by the natural logarithm, i.e., $\ln(n+1) \leq H(n) \leq \ln n + 1$. For unweighted cooperative games, the price of anarchy is $H(n)$ [Albers 09, Epstein et al. 09].

In this paper we study whether an optimal solution—which is a minimum-cost network establishing the required connections—represents an attractive, relatively stable state that agents might want to purchase. If the n agents buy an optimal solution, what extra cost does any agent incur compared to a strategy deviation? Such an optimal solution could be signaled by a central authority. The agents then adhere to this solution if the incentive of deviating is not too high. Moreover, an optimal solution is attractive for the following two reasons. (1) In Nash equilibria, there exist agents that pay a high cost compared to the average agent cost in an optimal solution. In a worst-case equilibrium, this cost factor can be as high as n ; even in a best-case equilibrium, the factor can be $H(n)$. With this information in mind, the agents might be interested in purchasing an optimal solution if the incentive of a strategy deviation is not too high. (2) The only known protocol to reach a good equilibrium, attaining a price of stability of $H(n)$, relies on optimal solutions. In [Anshelevich et al. 08b], it is shown that if the agents start in an optimal solution, then a sequence of improving moves converges to a Nash equilibrium whose cost is at most $H(n)$ times the optimum cost. Hence, if agents start in an optimal solution, they might as well consider remaining in that solution if the state has favorable properties.

Given this motivation, our work has relations to *signaling*, a basic concept in game theory. In signaling, a more-informed agent transmits information to less-informed agents so as to improve the outcome of the game. We refer the reader to [Gibbons 92, Spence 73, Spence 02] for some foundational work on signaling, where the last two articles address in particular signaling in economic markets. Recently, papers [Bro Miltersen and Sheffet 12, Emek et al. 12] explored signaling schemes for revenue maximization in auctions.

We evaluate the quality of optimal solutions by studying *approximate Nash equilibria* in which the equilibrium constraint is relaxed [Anshelevich et al. 08a, Chen and Roughgarden 10]. In a noncooperative game, a combination of strategies forms an α -approximate Nash equilibrium, for some $\alpha \geq 1$, if no agent can improve its cost by a factor of more than α , assuming that all the other agents adhere to their strategies. In cooperative games, a combination of strategies is an α -approximate strong Nash equilibrium if no coalition of agents can change strategy such that all agents of the coalition improve their cost by a factor of more than α .

We study the performance of optimal solutions for a variety of settings. The main conclusion is that optimal solutions represent good states for single-source games in undirected graphs. This holds for unweighted games, considering both

noncooperative and cooperative agent behavior, as well as for weighted games. On the other hand, in general source–sink games and in directed graphs, optimal solutions do not represent satisfying states.

1.1. Previous Work

Research on the network design game defined above was initiated in [Anshelevich et al. 08a]. In this first paper, the authors considered general cost-sharing schemes that are not restricted to Shapley cost-sharing. The cost of an edge may be split in an arbitrary way among agents. They considered undirected graphs. First, they studied single-source games in which each agent i has to connect one terminal to a common source. They showed that the cost of an optimal solution can be shared among the agents such that the resulting strategies form a Nash equilibrium. They also studied general source–sink games whereby each agent has to connect an arbitrary set of terminals. Here the cost of an optimal solution can be shared such that the agents' strategies form a 3-approximate Nash equilibrium.

In a second paper [Anshelevich et al. 08b], network design with Shapley cost-sharing was investigated. The authors first focused on unweighted games and showed that in directed and undirected graphs, the price of anarchy is n , while the price of stability is bounded above by $H(n)$. This upper bound of $H(n)$ is tight for directed graphs. Additionally, they studied weighted games and proved a lower bound of $\Omega(\max\{n, \log W\})$ on the price of stability, where W is the total weight of all the agents.

Further work on unweighted games was presented in [Chekuri et al. 07, Fiat et al. 06]. Both papers address single-source games in undirected graphs. The first showed that the price of anarchy is $O(\sqrt{n} \log^2 n)$ if agents join the game sequentially and perform best-response moves. In the second, it was proved that the price of stability is $O(\log \log n)$ if each vertex of the graph is the terminal of some agent. Weighted games in directed graphs were further investigated in [Chen and Roughgarden 10]. The authors of that work assumed that each agent has to connect a terminal pair (s_i, t_i) and proved that for every $\alpha = \Omega(\log w_{\max})$, the price of stability of $O(\alpha)$ -approximate Nash equilibria is $O((\log W)/\alpha)$. Here w_{\max} is the maximum weight of any agent. This trade-off also holds for undirected graphs.

Cooperative network design games were studied in [Aumann 59, Epstein et al. 09], where it was shown that the price of anarchy drops to $H(n)$. Finally, approximate pure Nash equilibria for a different class of graphical games were studied in [Nguyen and Tardos 09].

1.2. Our Contribution

We evaluate the stability of optimal solutions in network design games with Shapley cost-sharing, complementing the existing results for this classical cost-sharing mechanism. In Section 3, we present a comprehensive study of unweighted games. We focus mostly on single-source games in undirected graphs. First, for non-cooperative games, we prove that every optimal solution represents an $H(n)$ -approximate Nash equilibrium. We show that this bound is tight. There exist games in which an optimal solution does not form an α -approximate Nash equilibrium for $\alpha < H(n)$.

Then we investigate cooperative games in which agents may coordinate their actions. We consider a general scenario in which coalitions of up to c agents may be formed, for any $1 \leq c \leq n$. We prove that every optimal solution is a $2c(\ln(n/c) + 2)$ -approximate strong Nash equilibrium. Furthermore, we give a nearly matching lower bound: there exist games in which an optimal solution does not represent an α -approximate strong Nash equilibrium for $\alpha < c' \ln(n/c')$, where $c' = \min\{c, \lfloor n/e \rfloor\}$. Hence, for $c < \lfloor n/e \rfloor$, the bound is $\alpha < c \ln(n/c)$; for large coalitions of size $c \geq \lfloor n/e \rfloor$, the bound is $\lfloor n/e \rfloor$ and hence linear in n . This behavior is consistent with our upper bound.

Moreover, we consider general source-sink games, in which each agent has to connect an individual set of terminals, as well as directed graphs. In both cases, we show negative results, even for noncooperative games. There are general source-sink games for which an optimal solution is an $\Omega(n)$ -approximate Nash equilibrium. In directed graphs, the approximation guarantees that α can even be unbounded.

In Section 4, we study weighted games. We consider single-source games in undirected graphs. We show that in noncooperative games, every optimal solution is an α -approximate Nash equilibrium, where

$$\alpha = w_{\max} \sum_{k=0}^{n-1} \frac{1}{w_{\max} + k}.$$

This bound is again tight. Optimal solutions generally do not form α -approximate Nash equilibria for

$$\alpha < w_{\max} \sum_{k=0}^{n-1} \frac{1}{w_{\max} + k}.$$

The latter expression is bounded above by $w_{\max}(\ln(W/w_{\max}) + 1)$. Here w_{\max} and W denote again the maximum and total weight of the agents.

As for analysis techniques, our upper-bound results do not rely on potential functions. Instead we present direct, combinatorial analyses of the paths

connecting terminals in an optimal solution. For noncooperative games, the analysis is a bit involved. More specifically, we show that given any tree establishing the required connections and any coalition I of agents, there exists an agent $i \in I$ whose cost shares along the path from t_i to the source do not grow too much. Hence, for this agent, the incentive to perform a strategy change is not sufficiently high.

Here we finally relate our results to those of [Chen and Roughgarden 10] mentioned above for noncooperative games. In this paper, we evaluate the quality of optimal solutions, which are solutions of specific interest, and develop explicit bounds not resorting to O -notation. On the other hand, Chen and Roughgarden develop asymptotic trade-offs. For unweighted games, these trade-offs imply the existence of an $O(\log n)$ -approximate Nash equilibrium whose cost is within a constant factor of the optimum cost. The protocol starts in an optimal solution and then performs a sequence of improving deviations. Our results show that the protocol can, and indeed will, remain in the optimal solution.

2. Preliminaries

We formally define the network design games and game-theoretic concepts studied in this paper.

In *network design with Shapley cost-sharing* we are given a graph $G = (V, E, c)$ with a nonnegative cost function $c : E \mapsto \mathbb{R}_+^0$. As mentioned before, the graph may be directed or undirected. Associated with the graph are n selfish agents, whereby each agent i has to connect an individual set $T_i \subseteq V$ of terminals, $1 \leq i \leq n$. In a *single-source game*, each agent i has to connect one terminal t_i to a common source s , i.e., $T_i = \{t_i, s\}$. In a general *source-sink game*, T_i may be an arbitrary vertex set.

A strategy $S_i \subseteq E$ of an agent i is an edge set connecting the desired terminals. Considering all agents, we obtain a combination $\mathcal{S} = (S_1, \dots, S_n)$ of strategies. The cost of the edges used by the agents has to be fully covered. Shapley cost-sharing divides this cost in a “fair” manner. In an *unweighted game*, if k agents use an edge e , then each of them pays a share of $c(e)/k$. Thus, given a combination \mathcal{S} of strategies, the cost of an agent i , $1 \leq i \leq n$, is

$$\text{cost}_i(\mathcal{S}) = \sum_{e \in S_i} \frac{c(e)}{|\{j : e \in S_j\}|}.$$

In a *weighted game*, each agent i has a positive weight w_i and pays a share proportional to its weight. For each edge $e \in S_i$, agent i pays a share of $c(e)w_i/W_e$,

where $W_e = \sum_{j:e \in S_j} w_j$ is the total weight of the agents j using e in their strategies. The cost of agent i is $\text{cost}_i(\mathcal{S}) = \sum_{e \in S_i} c(e)w_i/W_e$.

We investigate the performance of optimal solutions using the concept of α -approximate Nash equilibria, which generalize standard Nash equilibria. Loosely speaking, in an α -approximate stable state, agents cannot improve their cost by a factor of more than α . More specifically, let $\alpha \geq 1$. In a noncooperative game, a combination \mathcal{S} of strategies forms an α -approximate Nash equilibrium if no agent i has a strategy change S'_i such that $\text{cost}_i(S_1, \dots, S'_i, \dots, S_n) < \text{cost}_i(\mathcal{S})/\alpha$, i.e., the cost after a strategy change is not smaller than $1/\alpha$ times the original cost. In cooperative games, a combination \mathcal{S} of strategies is an α -approximate strong Nash equilibrium if for every coalition of agents and associated strategy change, there exists one agent in the coalition whose cost does not decrease by a factor of more than α . Formally, for no nonempty coalition I of agents does there exist a strategy change S'_I such that $\text{cost}_i(S'_I, \mathcal{S}_{-I}) < \text{cost}_i(\mathcal{S})/\alpha$ holds for all $i \in I$. Here \mathcal{S}_{-I} is the vector of the original strategies of agents $i \notin I$.

3. Unweighted Games

In this section we study unweighted network design games with Shapley cost-sharing. We first consider the standard setting in which agents are noncooperating entities. Then we consider the setting in which agents cooperate and may form coalitions. For both scenarios, we focus on single-source games in undirected graphs. Finally, we address general source–sink games and games in directed graphs.

3.1. Noncooperative Games

We first prove an upper bound on the quality of optimal solutions and then give a matching lower bound.

Theorem 3.1. *In single-source games, every optimal solution represents an $H(n)$ -approximate Nash equilibrium.*

Proof. Let E_{opt} be the edge set used by an optimal solution to establish the required connections. Since we are studying single-source games, E_{opt} forms a tree. Consider the combination \mathcal{S} of strategies in which every agent i connects its terminal t_i to the common source s using only edges of E_{opt} . Let P_i be the simple path used by agent i , and let $\text{cost}_i(P_i)$ denote the corresponding cost paid by i within \mathcal{S}_{opt} . We observe that the path P_i is unique in E_{opt} . Furthermore,

if for any two agents i and i' , the paths P_i and $P_{i'}$ share a common vertex v , then the subpath from v to s in P_i is identical to the subpath from v to s in $P_{i'}$. This holds because otherwise, E_{opt} would contain cycles (since both P_i and $P_{i'}$ connect to s), contradicting the fact that E_{opt} is in fact a tree.

Now suppose that an agent i changes strategy and selects a different path Q_i , $Q_i \neq P_i$, in order to connect t_i to s . Let $\text{cost}_i(Q_i)$ be the associated cost incurred by agent i in performing this strategy change. We will show that

$$\text{cost}_i(P_i) \leq H(n)\text{cost}_i(Q_i), \tag{3.1}$$

which establishes the theorem.

Let $v_1, \dots, v_\ell, \ell \geq 2$, be the sequence of vertices whereby P_i and Q_i separate and merge again; Figure 1 presents an example. More specifically, starting at t_i , paths P_i and Q_i first traverse a common subpath $P_i(1) = Q_i(1)$ until they reach vertex v_1 , where the two paths separate. Vertex v_1 may be equal to t_i , in which case paths $P_i(1) = Q_i(1)$ are empty. After v_1 , path P_i traverses a subpath $P_i(2)$ while Q_i uses a subpath $Q_i(2)$. These subpaths use disjoint edge sets and meet again only at vertex v_2 . In general, suppose that P_i and Q_i merge at a vertex v_j , with j being even. Then P_i and Q_i traverse a common subpath $P_i(j+1) = Q_i(j+1)$ until they reach v_{j+1} , where P_i and Q_i separate into disjoint subpaths $P_i(j+2)$ and $Q_i(j+2)$, meeting again at v_{j+2} . Finally, let $P_i(\ell+1) = Q_i(\ell+1)$ be the subpath between v_ℓ and s . For an odd number j , the subpath $P_i(j) = Q_i(j)$ may be empty, in which case $v_{j-1} = v_j$ if $2 < j < \ell - 1$.

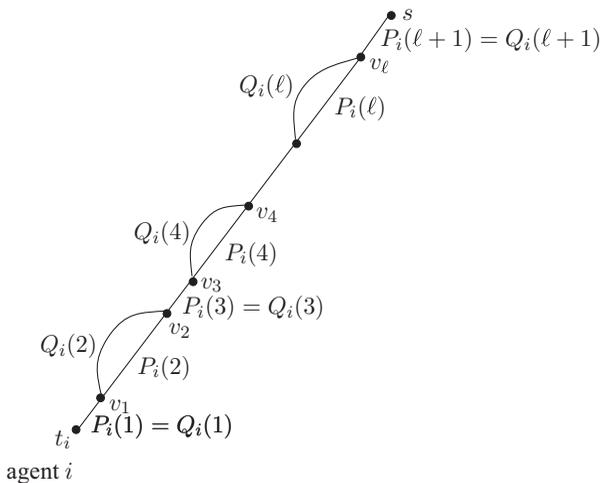


Figure 1. Paths P_i and Q_i .

Furthermore, for every even j , the subpath $Q_i(j)$ contains at least one edge that does not belong to E_{opt} , because the optimal solution does not contain cycles. Let $Q'_i(j)$, with $Q'_i(j) \subseteq Q_i(j)$, be the set of edges not contained in E_{opt} .

Let $\text{cost}_i(P_i(j))$ and $\text{cost}_i(Q_i(j))$ denote the costs paid by agent i on $P_i(j)$ and $Q_i(j)$, respectively, $1 \leq j \leq \ell + 1$. We have

$$\text{cost}_i(P_i) = \sum_{j=1}^{\ell+1} \text{cost}_i(P_i(j)) \quad \text{and} \quad \text{cost}_i(Q_i) = \sum_{j=1}^{\ell+1} \text{cost}_i(Q_i(j)),$$

where $\text{cost}_i(P_i(j)) = \text{cost}_i(Q_i(j))$ for every odd index j . We will prove that

$$\text{cost}_i(P_i(j)) \leq H(n)\text{cost}_i(Q_i(j))$$

for every even j , which implies inequality (3.1).

Consider a fixed even j and partition $P_i(j)$ into a sequence of maximal subpaths $P_i(j, 1), \dots, P_i(j, \ell_j)$ such that for every $1 \leq k \leq \ell_j$, the number of agents using a given edge e of $P_i(j, k)$ in E_{opt} is the same for all the edges of this subpath; cf. Figure 2. Such a partitioning is possible because, as argued at the beginning of this proof, if the path $P_{i'}$ of an agent i' meets P_i at some vertex, then the further subpaths of P_i and $P_{i'}$ in the direction to s are identical. Hence if an agent i' uses an edge of $P_i(j, k)$, then it uses all further edges of $P_i(j, k)$ and all the edges of $P_i(j, k + 1), \dots, P_i(j, \ell_j)$ in order to reach source s . In fact, agent i' also uses all edges of $P_i(j, k)$, by the definition of the subpaths $P_i(j, 1), \dots, P_i(j, \ell_j)$.

So let $n_i(j, k)$ be the number of agents using the edges of $P_i(j, k)$ within E_{opt} , for $1 \leq k \leq \ell_j$. Since the subpaths are maximal, $n_i(j, 1) < \dots < n_i(j, \ell_j)$. Since E_{opt} is a minimum-cost tree, we have $\text{cost}(P_i(j, k)) \leq \text{cost}(Q'_i(j))$, where

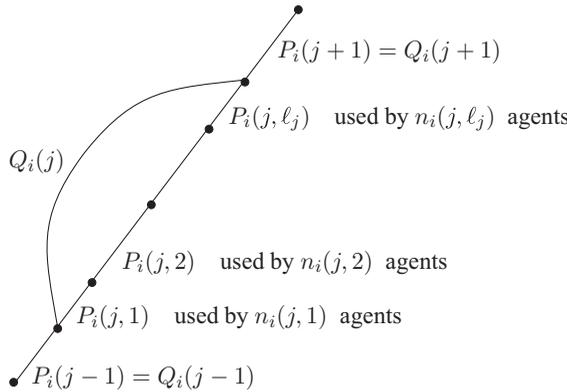


Figure 2. Subpaths $P_i(j)$ and $Q_i(j)$.

$\text{cost}(P_i(j, k))$ and $\text{cost}(Q'_i(j))$ denote the total edge costs of subpath $P_i(j, k)$ and edge set $Q'_i(j)$, respectively, $1 \leq k \leq \ell_j$. If we had $\text{cost}(P_i(j, k)) > \text{cost}(Q'_i(j))$, then in E_{opt} we could replace $P_i(j, k)$ by $Q'_i(j)$, obtaining a solution with a strictly smaller cost. The connectivity requirements would still be maintained, since agents using $P_i(j, k)$ in E_{opt} could traverse subpaths $P_i(j, k - 1), \dots, P_i(j, 1)$ and $Q_i(j)$ to reach v_j , from where they could again follow their original path to source s . In E_{opt} , agent i pays a share of $\text{cost}(P_i(j, k))/n_i(j, k)$ for $P_i(j, k)$, where $\text{cost}(P_i(j, k))$ is the total cost of the edges on $P_i(j, k)$. Summing over all k and making use of the fact that the sequence $n_i(j, k)$ is strictly increasing with $n_i(j, 1) \geq 1$, we obtain that the total cost paid by agent i on $P_i(j)$ is

$$\text{cost}_i(P_i(j)) = \sum_{k=1}^{\ell_j} \frac{\text{cost}(P_i(j, k))}{n_i(j, k)} \leq \sum_{k=1}^{\ell_j} \frac{\text{cost}(P_i(j, k))}{k} \leq H(n)\text{cost}(Q'_i(j)).$$

Since $Q'_i(j)$ is not part of E_{opt} , agent i has to cover the associated edge cost fully when traversing $Q_i(j)$, and hence $\text{cost}(Q'_i(j)) \leq \text{cost}_i(Q_i(j))$. We conclude that $\text{cost}_i(P_i(j)) \leq H(n)\text{cost}_i(Q_i(j))$. □

Theorem 3.2. *There exists a single-source game in which the unique optimal solution does not represent an α -approximate Nash equilibrium for any $\alpha < H(n)$.*

Proof. Consider a graph consisting of $n + 1$ vertices v_1, \dots, v_{n+1} and n edges $e_i = \{v_i, v_{i+1}\}$, $1 \leq i \leq n$; cf. Figure 3. Associated with each v_i , $1 \leq i \leq n$, is one agent that wishes to connect this vertex to the source $s = v_{n+1}$. Each edge e_i , $1 \leq i \leq n$, has cost 1. Additionally there is an edge $e_0 = \{v_1, v_{n+1}\}$ of cost $1 + \epsilon$, where $\epsilon > 0$ is an arbitrarily small constant. The unique optimal solution consists of the set of edges e_i , $1 \leq i \leq n$. In this solution, agent 1 pays a cost of $H(n)$. On the other hand, choosing edge e_0 , agent 1 incurs a cost of only $1 + \epsilon$. □

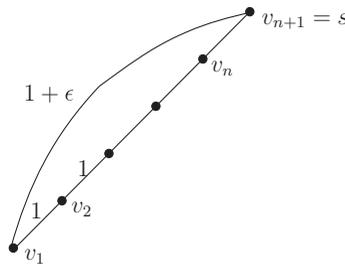


Figure 3. A single-source game without cooperation.

3.2. Cooperative Games

We study general cooperative games in which coalitions of up to c agents may be formed, for any $1 \leq c \leq n$.

Theorem 3.3. *In single-source games, every optimal solution represents an α -approximate strong Nash equilibrium, where $\alpha = 2c(\ln(n/c) + 2)$, if coalitions of size up to c are allowed.*

In order to establish the theorem, we first prove a property for trees T in which agents connect terminals to the root of T using the edges of the tree. The property holds for every tree T , but when we use the property in the proof of Theorem 3.3, T will be an optimal solution of a single-source game. So let T be an arbitrary tree with root s . There are n agents, each of which has to connect a terminal of T to s using the edges of T . Let A denote the set of agents i whose terminal t_i is different from s . For every agent $i \in A$, let P_i be the path from t_i to s in T . We partition P_i into maximal subpaths $P_i(1), \dots, P_i(l_i)$ such that for every subpath, the number of agents using the edges of the subpath does not vary. Let $n_i(j)$ be the number of agents using the edges of $P_i(j)$, $1 \leq j \leq l_i$. Define

$$N_i(T) = \sum_{j=1}^{l_i} \frac{1}{n_i(j)},$$

which intuitively is the sum of the fractions paid by agent i on $P_i(1), \dots, P_i(l_i)$, ignoring edge costs. Given these definitions, we present the following lemma, which states that in every nonempty coalition $I \subseteq A$, there exists an agent i whose value $N_i(T)$ is logarithmic in $|A|/|I|$.

Lemma 3.4. *Let T be an arbitrary tree and let A be the set of agents whose terminal is not equal to the root of T . For every $I \subseteq A$, $I \neq \emptyset$, there exists $i \in I$ satisfying*

$$N_i(T) \leq 2 \ln \left(\frac{2|A|}{|I|} \right) + 1.$$

Proof. We prove a slightly stronger bound on $N_i(T)$. Given T and A , a vertex $v \neq s$ in T is called a *branching vertex* if v has at least two children rooting subtrees both of which contain terminals. Let B be the set of branching vertices. We will prove

$$N_i(T) \leq 2 \ln \left(\frac{|A| + |B|}{|I|} \right) + 1. \quad (3.2)$$

The lemma then follows because $|B| < |A|$.

We prove (3.2) by induction on the number m of the edges of the tree T . In the base case we have $m = 1$. Tree T consists of a single edge $\{v, s\}$, and A is the set of agents that have to connect v to s . For every $I \subseteq A$, $I \neq \emptyset$, and $i \in I$ we have

$$N_i(T) = \frac{1}{|A|} \leq 1 \leq 2 \ln \left(\frac{|A| + |B|}{|I|} \right) + 1.$$

The last inequality holds because $|A| \geq |I|$ and $|B| = 0$.

Assume that (3.2) holds for trees consisting of up to $m - 1$ edges. We consider a tree T with m edges, where $m > 1$. If there is an agent $i \in I$ whose terminal t_i is equal to a child of s , then the analysis is simple. For this agent we have $N_i(T) \leq 1$, and as above, we conclude that

$$N_i(T) \leq 2 \ln \left(\frac{|A| + |B|}{|I|} \right) + 1,$$

because $|A| \geq |I|$ and $|B| \geq 0$. In the following, we assume that for no agent $i \in I$ is the terminal t_i equal to a child of s . We distinguish two cases depending on whether s has degree 1 or degree greater than 1.

We first study the case that s has degree 1. Let $\{s', s\}$ be the edge adjacent to s in T , and let T' be the tree rooted at s' . Let $A' \subseteq A$ be the set of agents i whose terminal t_i is a vertex of T' but not equal to the root s' . Then $I \subseteq A'$, because we assume that for no agent of I is the terminal equal to a child of s . For every $i \in I$, consider the path P_i from t_i to s and the path P'_i from t_i to s' . Obviously, P_i consists of P'_i followed by edge $\{s', s\}$. Partition both P_i and P'_i into maximal subpaths $P_i(1), \dots, P_i(l_i)$ and $P'_i(1), \dots, P'_i(l'_i)$ such that the edges of a subpath are used by a nonvarying number of agents. Let $n_i(j)$ and $n'_i(j)$ be the numbers of agents using $P_i(j)$ and $P'_i(j)$. We have $P_i(j) = P'_i(j)$ and hence $n_i(j) = n'_i(j)$, for $j = 1, \dots, l'_i - 1$. If the number $n'_i(l'_i)$ of agents using $P'_i(l'_i)$ is equal to the number of agents using edge $\{s', s\}$, then $l_i = l'_i$, and $P_i(l_i)$ consists of $P'_i(l_i)$ followed by $\{s', s\}$. Otherwise, $l_i = l'_i + 1$ as well as $P_i(l'_i) = P'_i(l'_i)$ and $P_i(l_i) = \{s', s\}$.

By the induction hypothesis, there exists an agent $i \in I$ satisfying

$$N_i(T') \leq 2 \ln \left(\frac{|A'| + |B'|}{|I|} \right) + 1,$$

where B' is the set of branching vertices in T' . In the following, we consider this fixed agent i . If $n'_i(l'_i)$ is equal to the number of agents using $\{s', s\}$, then we are done: as argued in the previous paragraph, $n_i(j) = n'_i(j)$, for $j = 1, \dots, l'_i - 1$,

and $l_i = l'_i$, which implies $n_i(l_i) = n'_i(l'_i)$. Hence

$$N_i(T) = N_i(T') \leq 2 \ln \left(\frac{|A'| + |B'|}{|I|} \right) + 1 \leq 2 \ln \left(\frac{|A| + |B|}{|I|} \right) + 1,$$

because $|A'| \leq |A|$ and $|B'| \leq |B|$.

If, on the other hand, $n'_i(l'_i)$ is not equal to the number of agents using $\{s', s\}$, then (a) there exists an agent in A whose terminal is equal to s' or (b) s' is a branching vertex. In case (a), we have $|A| > |A'|$, and in case (b), we have $|B| > |B'|$. Hence in both cases, $|A| + |B| > |A'| + |B'|$. Again $n_i(j) = n'_i(j)$, for $j = 1, \dots, l'_i - 1$. Since $l_i = l'_i + 1$ and $P_i(l'_i) = P'_i(l'_i)$, we have $n_i(l'_i) = n'_i(l'_i)$ and $n_i(l_i) = 1/|A|$, because edge $\{s', s\}$ is used by all the agents of A . We obtain

$$\begin{aligned} N_i(T) &= N_i(T') + \frac{1}{|A|} \leq 2 \ln \left(\frac{|A'| + |B'|}{|I|} \right) + 1 + \frac{2}{2|A|} \\ &\leq 2 \left(\ln \left(\frac{|A'| + |B'|}{|I|} \right) + \frac{1}{|A| + |B|} \right) + 1 \\ &\leq 2 \left(\ln(|A'| + |B'|) + \frac{1}{|A'| + |B'| + 1} - \ln(|I|) \right) + 1. \end{aligned}$$

The second inequality holds because $|A| > |B|$ and hence $2|A| > |A| + |B|$. The third inequality follows since $|A| + |B| \geq |A'| + |B'| + 1$. For every positive integer K , we have $\ln K + 1/(K + 1) \leq \ln(K + 1)$. Setting $K = |A'| + |B'|$ and observing again that $|A| + |B| \geq |A'| + |B'| + 1$, we obtain as desired

$$N_i(T) \leq 2(\ln(|A'| + |B'| + 1) - \ln(|I|)) + 1 \leq 2 \ln \left(\frac{|A| + |B|}{|I|} \right) + 1.$$

It remains to study the case that the root s of T has degree greater than 1. Hence there are at least two subtrees linked to s . We divide T into two trees T_1 and T_2 . Tree T_1 is obtained from T by removing an arbitrary subtree T' rooted at a child s' of s . The edge $\{s', s\}$ is removed as well. Tree T' is attached below a new root vertex s_2 using an edge $\{s', s_2\}$ and, together with s_2 and $\{s', s_2\}$, forms the second tree T_2 . Intuitively, T is obtained from T_1 and T_2 by simply merging the roots of T_1 and T_2 .

Recall that A and I do not contain agents whose terminal is equal to the root s of T . Let $A_j \subseteq A$, for $j \in \{1, 2\}$, be the set of agents i whose terminal t_i is in T_j . Similarly, let $I_j \subseteq I$, for $j \in \{1, 2\}$, be the set of agents whose terminal is in T_j . Since T_1 and T_2 contain different subtrees of T , sets A_1 and A_2 as well as I_1 and I_2 are disjoint. Hence $|A| = |A_1| + |A_2|$ and $|I| = |I_1| + |I_2|$. In the tree T , agents of the set A_1 do not share edges with agents of A_2 when connecting their terminals to s . Hence, for every agent $i \in I_1$, the structure of P_i in T is the same as in T_1 . Again P_i is the path connecting t_i to the root of the tree. Analogously,

for every agent $i \in I_2$, the structure of P_i in T is the same as in T_2 , because in T_2 we added an edge $\{s', s_2\}$ to a new root s_2 simulating the edge $\{s', s\}$ in T . This implies $N_i(T) = N_i(T_1)$ for every $i \in I_1$, and $N_i(T) = N_i(T_2)$ for every $i \in I_2$.

If one of the two sets I_1, I_2 is empty, we are done: Assume without loss of generality that $I_2 = \emptyset$. By the induction hypothesis, there exists an agent $i \in I_1$ with

$$N_i(T_1) \leq 2 \ln \left(\frac{|A_1| + |B_1|}{|I_1|} \right) + 1 = 2 \ln \left(\frac{|A_1| + |B_1|}{|I|} \right) + 1,$$

where B_1 is the set of branching vertices in T_1 . Since $|A_1| \leq |A|$ and $|B_1| \leq |B|$, we obtain

$$N_i(T) = N_i(T_1) \leq 2 \ln \left(\frac{|A_1| + |B_1|}{|I|} \right) + 1 \leq 2 \ln \left(\frac{|A| + |B|}{|I|} \right) + 1.$$

If both I_1 and I_2 are not empty, then we consider the smaller of the two ratios

$$\frac{|A_j| + |B_j|}{|I_j|},$$

where $j = 1, 2$. Assume without loss of generality that

$$\frac{|A_1| + |B_1|}{|I_1|} \leq \frac{|A_2| + |B_2|}{|I_2|}.$$

Then

$$\frac{|A_1| + |B_1|}{|I_1|} \leq \frac{|A_1| + |B_1| + |A_2| + |B_2|}{|I_1| + |I_2|} = \frac{|A| + |B|}{|I|}.$$

By the induction hypothesis, there exists again an agent $i \in I_1$ with

$$N_i(T_1) \leq 2 \ln \left(\frac{|A_1| + |B_1|}{|I_1|} \right) + 1.$$

We conclude that for this agent i ,

$$N_i(T) = N_i(T_1) \leq 2 \ln \left(\frac{|A_1| + |B_1|}{|I_1|} \right) + 1 \leq 2 \ln \left(\frac{|A| + |B|}{|I|} \right) + 1.$$

□

Proof of Theorem 3.3. Consider any optimal solution and let E_{opt} be the corresponding edge set. Moreover, let \mathcal{S} be the combination of strategies in which every agent i connects its terminal t_i to the common source s using only edges of E_{opt} . In order to prove the theorem, we show that if any nonempty coalition I of at most c agents changes strategy, then there exists an agent $i \in I$ whose cost before and after a strategy change satisfies

$$\frac{1}{\alpha} \text{cost}_i(\mathcal{S}) \leq \text{cost}_i(\mathcal{S}_I, \mathcal{S}_{-I}), \tag{3.3}$$

where $\alpha = 2c(\ln(n/c) + 2)$.

If a coalition I contains an agent i whose terminal t_i is equal to the source s , then there is nothing to show, because for this agent, $\text{cost}_i(\mathcal{S}) = 0$ and (3.3) trivially holds. Therefore, in the following we always consider nonempty coalitions I not containing an agent i whose terminal is equal to s .

Let A be the set of agents whose terminal is not equal to s . Consider any nonempty coalition $I \subseteq A$ of size at most c . The optimal solution E_{opt} forms a tree, and hence by Lemma 3.4, there exists an agent $i \in I$ with

$$N_i(E_{\text{opt}}) \leq 2 \ln \left(\frac{2|A|}{|I|} \right) + 1.$$

Fix this agent i . We will prove that if I performs any strategy change, then for this agent i , inequality (3.3) holds.

For agent i , let P_i be the path connecting t_i to s in E_{opt} . Let Q_i be the path used by the agent when I changes strategy. As in the proof of Theorem 3.1, we partition P_i and Q_i into subpaths $P_i(1), \dots, P_i(l+1)$ and $Q_i(1), \dots, Q_i(l+1)$ along the vertices v_1, \dots, v_l where P_i and Q_i separate and merge. Let $\text{cost}_i(P(j))$ be the cost incurred by agent i for $P_i(j)$ before the strategy change, $1 \leq j \leq l+1$. Similarly, let $\text{cost}_i(Q(j))$ be the cost paid by the agent for $Q_i(j)$ after the strategy change, $1 \leq j \leq l+1$. For every odd number j , we have $P_i(j) = Q_i(j)$ and hence

$$\frac{1}{|I|} \text{cost}_i(P(j)) \leq \text{cost}_i(Q(j)),$$

because at most $|I| - 1$ additional agents of I can join edges of $P_i(j)$ after the strategy change. Since $|I| \leq c$, this implies

$$\frac{1}{\alpha} \text{cost}_i(P(j)) \leq \text{cost}_i(Q(j))$$

for every odd number j , $1 \leq j \leq l+1$. In the following, we show that the last inequality also holds for every even number j .

For every even j we partition $P_i(j)$ into maximal subpaths $P_i(j, 1), \dots, P_i(j, l_j)$ such that all the edges of a subpath $P_i(j, k)$ are used by the same number $n_i(j, k)$ of agents, $1 \leq k \leq l_j$, considering the time before a strategy change. Let $Q'_i(j) \subseteq Q_i(j)$ be the nonempty set of edges not contained in E_{opt} . For every path π , let $\text{cost}(\pi)$ be the total cost of the edges of π . Then $\text{cost}(P_i(j, k)) \leq \text{cost}(Q'_i(j))$ for $1 \leq k \leq l_j$, and $\text{cost}(Q'_i(j)) \leq \text{cost}(Q_i(j))$. Hence

$$\text{cost}_i(P_i(j)) = \sum_{k=1}^{l_j} \frac{\text{cost}(P_i(j, k))}{n_i(j, k)} \leq \text{cost}(Q'_i(j)) \sum_{k=1}^{l_j} \frac{1}{n_i(j, k)}.$$

Consider the partitioning of P_i into maximal subpaths such that the edges of a subpath are used by the same number of agents. Each of $P_i(j, 1), \dots, P_i(j, l_j)$

occurs as a subpath or as a portion of a subpath in this partitioning, and hence $\sum_{k=1}^j 1/n_i(j, k) \leq N_i(E_{\text{opt}})$. Moreover, $\text{cost}_i(Q'_i(j)) \geq \text{cost}(Q'_i(j))/|I|$, because the cost of the edges of $Q'_i(j)$, which are not part of E_{opt} , must be fully covered by the coalition I , and agent i pays a share of at least $1/|I|$. Thus

$$\text{cost}_i(P_i(j)) \leq |I| \text{cost}_i(Q'_i(j)) N_i(E_{\text{opt}}) \leq |I| \cdot N_i(E_{\text{opt}}) \text{cost}_i(Q_i(j)).$$

By our choice of agent i and Lemma 3.4,

$$N_i(E_{\text{opt}}) \leq 2 \ln \left(\frac{2|A|}{|I|} \right) + 1 \leq 2 \ln \left(\frac{2n}{|I|} \right) + 1.$$

We obtain

$$\begin{aligned} \text{cost}_i(P_i(j)) &\leq |I| \left(2 \ln \left(\frac{2n}{|I|} \right) + 1 \right) \text{cost}_i(Q_i(j)) \\ &\leq c \left(2 \ln \left(\frac{2n}{c} \right) + 1 \right) \text{cost}_i(Q_i(j)). \end{aligned}$$

The last inequality holds because $|I|(2 \ln(2n/|I|) + 1)$ is increasing in $|I|$. We conclude that

$$\text{cost}_i(P_i(j)) \leq 2c \left(\ln \left(\frac{n}{c} \right) + 2 \right) \text{cost}_i(Q_i(j)),$$

and, as desired,

$$\frac{1}{\alpha} \text{cost}_i(P_i(j)) \leq \text{cost}_i(Q_i(j)).$$

□

Theorem 3.5. *There exists a single-source game allowing coalitions of size up to c in which the unique optimal solution does not represent an α -approximate strong Nash equilibrium for any $\alpha < c' \ln(n/c')$, where $c' = \min\{c, \lfloor n/e \rfloor\}$.*

Proof. Consider the graph depicted in Figure 4, which is a modification of the graph used in the proof of Theorem 3.2. Set $c' = \min\{c, \lfloor n/e \rfloor\}$. There are $n - c' + 2$ vertices $v_{c'}, v_{c'+1}, \dots, v_{n+1}$ that are connected by edges $e_i = \{v_i, v_{i+1}\}$, for $c' \leq i \leq n$. Associated with vertex $v_{c'}$ are c' agents that wish to connect this vertex to the source $s = v_{n+1}$. Associated with each vertex v_i , $c' < i \leq n$, is one vertex having to connect v_i to s . Each edge $e_i = \{v_i, v_{i+1}\}$, $c' \leq i \leq n$, has a cost of 1. The graph contains an additional edge $e_0 = \{v_{c'}, v_{n+1}\}$ of cost $1 + \epsilon$. The unique optimal solution consists of edges e_i , $c' \leq i \leq n$. Consider the coalition I of the c' agents having to connect $v_{c'}$ to s . In the optimal solution, each of these agents incurs a cost of

$$\frac{1}{c'} + \frac{1}{c'+1} + \dots + \frac{1}{n} \geq \ln(n+1) - \ln c' \geq \ln \left(\frac{n}{c'} \right).$$

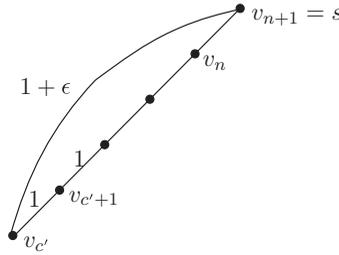


Figure 4. A single-source game with cooperation.

If I changes strategy and chooses e_0 to connect to s , then each of the agents incurs a cost of $(1 + \epsilon)/c'$ only. Hence the optimal solution does not represent an α -approximate strong Nash equilibrium for $\alpha < c' \ln(n/c')$. \square

3.3. Source–Sink Games and Directed Graphs

It is natural to ask whether the results of the previous sections can be extended (a) to general source–sink games in which each agent has to connect an individual set of terminals or (b) to directed graphs. Unfortunately, this is not the case. Even for noncooperative games, we can show high lower bounds on the approximation factor α .

Theorem 3.6. *There exists a general source–sink game in which the unique optimal solution represents an α -approximate Nash equilibrium with $\alpha = \Omega(n)$.*

Proof. Consider the graph depicted in Figure 5. There are n vertices v_1, \dots, v_n that are connected by edges $e_i = \{v_i, v_{i+1}\}$, $1 \leq i \leq n - 1$. Furthermore, there are vertices u_1, \dots, u_{n-1} , and w_1, \dots, w_{n-1} with corresponding edges $\{u_i, v_i\}$ and $\{v_{i+1}, w_i\}$, $1 \leq i \leq n - 1$. Agent i , $1 \leq i \leq n - 1$, has to connect u_i and w_i . There are two additional vertices v_0 and v_{n+1} with associated edges $e_0 = \{v_0, v_1\}$ and $e_n = \{v_n, v_{n+1}\}$. Agent n has to connect terminals v_0 and v_{n+1} . All edges mentioned so far have a cost of 1. Finally, there is an edge $e' = \{v_0, v_{n+1}\}$ of cost $2 + \epsilon$. The unique optimal solution purchases all the edges e_i , $0 \leq i \leq n$, in addition to $\{u_i, v_i\}$ and $\{v_{i+1}, w_i\}$, $1 \leq i \leq n - 1$. In this solution, agent n has to pay a cost of

$$2 + \frac{n - 1}{2} \geq \frac{n}{2},$$

whereas a purchase of edge e' incurs a cost of $2 + \epsilon$ only. \square

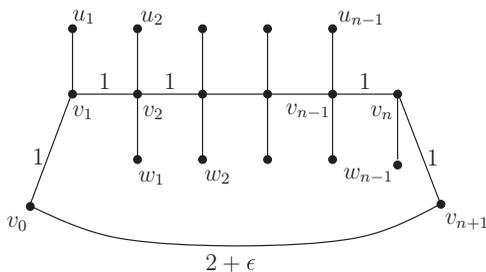


Figure 5. A source-sink game.

Theorem 3.7. *For every constant C , there exist single-source games in directed graphs in which an optimal solution does not form a C -approximate Nash equilibrium.*

Proof. Consider the graph used in Theorem 3.2. We now orient all edges toward the source s , i.e., we have directed edges (v_i, v_{i+1}) , $1 \leq i \leq n$, and (v_1, v_{n+1}) . The edge (v_1, v_{n+1}) adjacent to v_1 still has cost $1 + \epsilon$. On the other edges we modify the costs such that (v_2, v_3) costs $2C$ and the remaining edges (v_1, v_2) as well as (v_i, v_{i+1}) , for every $3 \leq i \leq n$, cost 0. The optimal solution still consists of edges (v_i, v_{i+1}) , $1 \leq i \leq n$. In this solution, agent 1 pays a cost of C , whereas a purchase of edge (v_1, v_{n+1}) reduces its cost to $1 + \epsilon$. \square

4. Weighted Games

In this section, we consider weighted network design games and generalize the results for noncooperative single-source games in undirected graphs developed in Section 3.1. More specifically, we extend Theorems 3.1 and 3.2 and give tight upper and lower bounds on the value of α such that every optimal solution represents an α -approximate Nash equilibrium.

We scale the agents' weights such that the minimum weight is equal to 1 and hence $w_i \geq 1$, for all the agents. Let $w_{\max} = \max_{1 \leq i \leq n} w_i$ be the maximum weight of all the agents. The expression

$$w_{\max} \sum_{k=0}^{n-1} \frac{1}{w_{\max} + k}$$

given in the following theorem is bounded above by $w_{\max}(\ln(W/w_{\max}) + 1)$.

Theorem 4.1. *In single-source games, every optimal solution represents an α -approximate Nash equilibrium, where $\alpha = w_{\max} \sum_{k=0}^{n-1} 1/(w_{\max} + k)$.*

Proof. The proof is very similar to that of Theorem 3.1. In every optimal solution E_{opt} , we consider an arbitrary agent i along with its routing path P_i in E_{opt} . Suppose that agent i changes strategy and selects a different path $Q_i \neq P_i$. Paths P_i and Q_i are partitioned in the same way as before into subpaths $P_i(j)$ and $Q_i(j)$, $1 \leq j \leq \ell + 1$. We show that for every even j , the costs incurred by agent i on $P_i(j)$ and $Q_i(j)$ satisfy

$$\text{cost}_i(P_i(j)) \leq \text{cost}_i(Q_i(j)) \cdot w_{\max} \sum_{k=0}^{n-1} \frac{1}{w_{\max} + k}, \quad (4.1)$$

which establishes the theorem, because for odd indices j , we have $P_i(j) = Q_i(j)$, and hence $\text{cost}_i(P_i(j)) = \text{cost}_i(Q_i(j))$.

To prove (4.1), for even j , we partition $P_i(j)$ again into subpaths $P_i(j, 1), \dots, P_i(j, \ell_j)$ of maximum length such that the number of agents using a given subpath does not change on the subpath. Let $W_i(j, k)$, $1 \leq k \leq \ell_j$, be the total weight of agents using the edges of $P_i(j, k)$. The cost incurred by agent i on $P_i(j)$ is

$$\text{cost}_i(P_i(j)) = \sum_{k=1}^{\ell_j} \text{cost}(P_i(j, k)) \frac{w_i}{W_i(j, k)} \leq \sum_{k=1}^{\ell_j} \text{cost}(Q'_i(j)) \frac{w_i}{W_i(j, k)},$$

where $Q'_i(j)$ is again the set of edges of $Q_i(j)$ not contained in E_{opt} . Hence

$$\text{cost}_i(P_i(j)) \leq \sum_{k=1}^{\ell_j} \text{cost}_i(Q_i(j)) \frac{w_i}{W_i(j, k)}.$$

Since the weight of each agent is at least 1, agent i uses path $P_i(j, 1)$ and the number of agents using subpaths $P_i(j, k)$ is strictly increasing as k increases, we have $W_i(j, k) \geq w_i + k - 1$. Thus

$$\text{cost}_i(P_i(j)) \leq \text{cost}_i(Q_i(j)) \sum_{k=1}^{\ell_j} \frac{w_i}{(w_i + k - 1)}.$$

The function $f(w) = \sum_{k=1}^{\ell_j} w/(w + k - 1)$ is increasing in w . We conclude that

$$\begin{aligned} \text{cost}_i(P_i(j)) &\leq \text{cost}_i(Q_i(j)) \cdot w_{\max} \sum_{k=1}^{\ell_j} \frac{1}{w_{\max} + k - 1} \\ &\leq \text{cost}_i(Q_i(j)) \cdot w_{\max} \sum_{k=0}^{n-1} \frac{1}{w_{\max} + k}. \end{aligned}$$

□

Theorem 4.2. *There exists a single-source game in which the unique optimal solution does not represent an α -approximate Nash equilibrium for any*

$$\alpha < w_{\max} \sum_{k=0}^{n-1} \frac{1}{w_{\max} + k}.$$

Proof. The game is identical to that used in the proof of Theorem 3.2, except that agents now have weights. Agent 1 associated with v_1 has weight $w_{\max} > 1$. All the other agents have weight 1. In the optimal solution, consisting of the edges $e_i = \{v_i, v_{i+1}\}$, $1 \leq i \leq n$, agent 1 incurs a cost of $w_{\max}/(w_{\max} + i - 1)$ for edge e_i . Summing over all i , we obtain a total cost of $w_{\max} \sum_{i=0}^{n-1} 1/(w_{\max} + i)$. If agent 1 changes strategy and purchases edge e_0 , its cost is $1 + \epsilon$. \square

5. Conclusions and Open Problems

In this paper we have studied a basic network design game and analyzed the stability of optimal solutions. For unweighted games in undirected graphs, every optimal solution represents an $H(n)$ -approximate Nash equilibrium, and this bound is tight. A major open problem is to find out whether one can achieve better guarantees by slightly increasing the cost of the solution: does there exist a solution whose cost is within a constant factor of the optimal cost and that forms an α -approximate Nash equilibrium for some $\alpha \in o(\log n)$? If the answer is positive, can such a solution be constructed explicitly, or is it attained via best-response moves from an optimal solution? At the time of this writing, the authors do not know how to resolve these questions.

More generally, it would be interesting to consider other games and study stability properties of optimal solutions.

Acknowledgments. We thank an anonymous referee for constructive comments that helped us to improve the presentation of this paper.

References

- [Albers 09] S. Albers. “On the Value of Coordination in Network Design.” *SIAM Journal on Computing* 38 (2009), 2273–2302.
- [Alon et al. 10] N. Alon, E. D. Demaine, M. Hajiaghayi, and T. Leighton. “Basic Network Creation Games.” In *Proc. 22nd Annual ACM Symposium on Parallel Algorithms and Architectures*, pp. 106–113, 2010.

- [Anshelevich et al. 08a] E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. “Near Optimal Network Design with Selfish Agents.” *Theory of Computing* 4 (2008), 77–109.
- [Anshelevich et al. 08b] E. Anshelevich, A. Dasgupta, J. M. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. “The Price of Stability for Network Design with Fair Cost Allocation.” *SIAM Journal on Computing* 38 (2008), 1602–1623.
- [Aumann 59] R. J. Aumann. “Acceptable Points in General Cooperative n -Person Games.” In *Contributions to the Theory of Games*, vol. IV, edited by A. W. Tucker and R. D. Luce, Annals of Mathematics Studies 40, pp. 287–324. Princeton University Press, 1959.
- [Bro Miltersen and Sheffet 12] P. Bro Miltersen and O. Sheffet. “Send Mixed Signals: Earn More, Work Less.” In *Proc. 13th ACM Conference on Electronic Commerce*, pp. 234–247, 2012.
- [Chekuri et al. 07] C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. Naor, and A. Orda. “Non-cooperative Multicast and Facility Location Games.” *IEEE Journal on Selected Areas in Communications* 25:6 (2007), 1193–1206.
- [Chen and Roughgarden 10] H.-L. Chen and T. Roughgarden. “Network Design with Weighted Players.” *Theory of Computing Systems* 45 (2009), 302–324.
- [Chen et al. 10] H.-L. Chen, T. Roughgarden, and G. Valiant. “Designing Network Protocols for Good Equilibria.” *SIAM Journal on Computing* 39 (2010), 1799–1832.
- [Corbo and Parkes 05] J. Corbo and D. Parkes. “The Price of Selfish Behavior in Bilateral Network Formation.” In *Proc. 24th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 99–107, 2005.
- [Demaine et al. 09] E. D. Demaine, M. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. “The Price of Anarchy in Cooperative Network Creation Games.” *SIGecom Exchanges* 8:2 (2009), 2.
- [Devanur et al. 05] N. R. Devanur, M. Mihail, and V. V. Vazirani. “Strategyproof Cost-Sharing Mechanisms for Set Cover and Facility Location Games.” *Decision Support Systems*, 39 (2005), 11–22.
- [Emek et al. 12] Y. Emek, M. Feldman, I. Gamzu, R. Paes Leme, and M. Tennenholtz. “Signaling Schemes for Revenue Maximization.” In *Proc. 13th ACM Conference on Electronic Commerce*, pp. 514–531, 2012.
- [Epstein et al. 09] A. Epstein, M. Feldman, and Y. Mansour. “Strong Equilibrium in Cost Sharing Connection Games.” *Games and Economic Behavior* 67:1 (2009), 51–68.
- [Fabrikant et al. 03] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker. “On a Network Creation Game.” In *Proc. 22nd Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 347–351, 2003.
- [Fiat et al. 06] A. Fiat, H. Kaplan, M. Levy, S. Olonetsky, and R. Shabo. “On the Price of Stability for Designing Undirected Networks with Fair Cost Allocations.” In *Proc. 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, Springer LNCS 4051, pp. 608–618, 2006.
- [Gibbons 92] R. Gibbons. *A Primer in Game Theory*. Prentice Hall, Harlow, 1992.

- [Gupta et al. 08] A. Gupta, A. Srinivasan, and E. Tardos. “Cost-Sharing Mechanisms for Network Design.” *Algorithmica* 50 (2008), 98–119.
- [Halevi and Mansour 07] Y. Halevi and Y. Mansour. “A Network Creation Game with Nonuniform Interests.” In *Proc. 3rd International Workshop on Internet and Network Economics (WINE)*, Springer LNCS 4858, pp. 287–292, 2007.
- [Hoefler 09] M. Hoefler. “Non-cooperative Tree Creation.” *Algorithmica* 53 (2009), 104–131.
- [Koutsoupias and Papadimitriou 09] E. Koutsoupias and C. Papadimitriou. “Worst-Case Equilibria.” *Computer Science Review* 3:2 (2009), 65–69.
- [Laoutaris et al. 08] N. Laoutaris, L. J. Poplawski, R. Rajaraman, R. Sundaram, and S.-H. Teng. “Bounded Budget Connection (BBC) Games or How to Make Friends and Influence People, on a Budget.” In *Proc. 27th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pp. 165–174, 2008.
- [Nguyen and Tardos 09] T. Nguyen and E. Tardos. “Approximate Pure Nash Equilibria via Lovász Local Lemma.” In *Proc. 5th International Workshop on Internet and Network Economics (WINE)*, LNCS 5929, pp. 160–171. Springer, 2009.
- [Pal and Tardos 03] M. Pal and E. Tardos. “Group Strategyproof Mechanisms via Primal–Dual Algorithms.” In *Proc. 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp. 584–593, 2003.
- [Spence 73] M. Spence. “Job Market Signaling.” *Quarterly Journal of Economics* 87:3 (1973), 355–374.
- [Spence 02] M. Spence. “Signaling in Retrospect and the Informational Structure of Markets.” *American Economic Review* 92:3 (2002), 434–425.
- [Vetta 02] A. Vetta. “Nash Equilibria in Competitive Societies with Applications to Facility Location, Traffic Routing and Auctions.” In *Proc. 43rd Annual IEEE Symposium on the Foundations of Computer Science*, pp. 416–425, 2002.

Susanne Albers, Department of Computer Science, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany (albers@informatik.hu-berlin.de)

Pascal Lenzen, Department of Computer Science, Humboldt-Universität zu Berlin, Unter den Linden 6, 10099 Berlin, Germany (lenzen@informatik.hu-berlin.de)