On the Value of Coordination in Network Design

Susanne Albers*

Abstract

We study network design games where n self-interested agents have to form a network by purchasing links from a given set of edges. We consider Shapley cost sharing mechanisms that split the cost of an edge in a fair manner among the agents using the edge. It is well known that the price of anarchy of these games is as high as n. Therefore, recent research has focused on evaluating the price of stability, i.e. the cost of the best Nash equilibrium relative to the social optimum.

In this paper we investigate to which extent coordination among agents can improve the quality of solutions. We resort to the concept of *strong Nash equilibria*, which were introduced by Aumann and are resilient to deviations by coalitions of agents. We analyze the price of anarchy of strong Nash equilibria and develop lower and upper bounds for unweighted and weighted games in both directed and undirected graphs. These bounds are tight or nearly tight for many scenarios. It shows that using coordination, the price of anarchy drops from linear to logarithmic bounds.

We complement these results by also proving the first super-constant lower bound on the price of stability of standard equilibria (without coordination) in undirected graphs. More specifically, we show a lower bound of $\Omega(\log W/\log\log W)$ for weighted games, where W is the total weight of all the agents. This almost matches the known upper bound of $O(\log W)$. Our results imply that, for most settings, the worst-case performance ratios of strong coordinated equilibria are essentially always as good as the performance ratios of the best equilibria achievable without coordination. These settings include unweighted games in directed graphs as well as weighted games in both directed and undirected graphs.

*Department of Computer Science, University of Freiburg, Georges Köhler Allee 79, 79110 Freibug, Germany. salbers@informatik.uni-freiburg.de Work supported by the German - Israeli Foundation for Scientific Research & Development, project G-783-61.6/2003 and by the German Research Foundation, project AL 464/4-2. This work was done while on sabbatical leave at Carnegie Mellon University, Pittsburgh, and supported by the ALADDIN project.

1 Introduction

Communication networks are pervasive and critical to modern society. Nonetheless, the formation and evolution of large networks is not well understood, a major reason being that these networks typically are not built by a central authority but rather by many economic agents that have selfish interests. For this reason, research on network design has focused on game-theoretic approaches over the past years, see e.g. [2, 3, 5, 6, 7, 8, 9, 11, 14, 22, 23].

We study network design games that have received a lot of attention recently [2, 3, 6, 7, 13, 16] and are simple, yet powerful enough to capture the two most important objectives of agents: connection establishment and cost minimization. Consider a directed or undirected graph G where each edge e has a non-negative cost $c(e)$. There are n agents, each of which has to connect a set of terminals. The agents form a network by selecting edges. A strategy S_i of an agent i is a set of edges connecting the desired terminals. The cost of the edges used by all the agents has to be covered. A fundamental cost sharing mechanism is *Shapley cost sharing*, which was proposed by Anshelevich et al. [3] for network design games and has been studied with respect to other networking problems as well [12, 15]. In Shapley cost sharing, the cost of an edge is shared in a fair manner among the agents using that edge. In an *unweighted game* if k agents use an edge e in their strategies, then each of these agents pays a share of $c(e)/k$. In a *weighted game*, each agent i has a weight w_i and contributes a share of $c(e)w_i/W_e$, where W_e is the total weight of agents using e. We are interested in stable networks where no agent has the incentive to deviate from its strategy. Stability is modeled by considering Nash equilibria. A combination $\mathcal{S} = (S_1, \ldots, S_n)$ of strategies forms a Nash equilibrium if no agent has a better strategy with a strictly smaller cost if all the other agents adhere to their strategies. A widely accepted performance measure to evaluate the quality of Nash equilibria is the *price of anarchy* [19], which is the maximum ratio of the total cost incurred by any Nash equilibrium to the cost spent by the social optimum. Unfortunately, for our network design games, the price of anarchy is as high as n. As an alternative quality measure, Anshelevich et al. [3] proposed the *price of stability* which is the ratio of the best Nash equilibrium relative to the social optimum. Anshelevich et al. [3] proved that the price of stability in unweighted network design games is $O(\log n)$.

The scenario described so far assumes that agents are completely non-cooperative, isolated entities. However for long-term decisions such as network design, given today's communication infrastructure, this assumption is not entirely realistic. It is more natural that agents will discuss possible strategies and, as in other economic markets, form coalitions taking strategic actions that are beneficial to all members of the group. In such cooperative environments we seek again stable solutions. In this context, Aumann [4] in 1959 introduced the concept of *strong Nash equilibria*, which ensure stability against deviations by every conceivable coalition of agents. More specifically, no coalition can cooperatively deviate in a way that benefits all its members, taking the actions of the agents outside the coalition as given. With respect to network design, an important question is if coordination among agents yields strictly better solutions. Is it possible to achieve significant improvements? We prove that this is the case. When coordination is allowed, the price of anarchy of strong Nash equilibria drops from n to $O(\log n)$ in unweighted games. Similar improvements show in weighted games. Obviously, any strong Nash equilibrium is a standard Nash equilibrium, which is immune to deviations of single agents. Hence strong Nash equilibria cannot be better than the best standard Nash equilibria. A second natural question is how strong Nash equilibria rank relative to the best standard Nash equilibrium. When coordination is allowed, is the worst-case performance of stable states close to that of the best stable states achievable without cooperation? We answer this question in the affirmative in terms of anarchy and stability measures. For most settings, the price of anarchy of strong Nash equilibria is essentially always as good as the corresponding stability bounds of standard equilibria. These settings include unweighted games in directed graphs as well as weighted games in both directed and undirected graphs.

Previous results: Research on the network design games defined above was initiated by Anshelevich et al. [2]. In this first paper the authors considered general cost sharing schemes that are not restricted to Shapley mechanisms. Anshelevich et al. studied undirected graphs and first addressed scenarios where each agent has to connect one terminal to a common destination. They designed Nash equilibria whose cost is equal to the cost of the optimum. Furthermore Ashelevich et al. [2] investigated the general scenario that each agent has to connect a set of terminals. In this case there are graphs that do not admit Nash equilibria. The authors therefore studied *α-approximate Nash equilibria* in which no agent can improve its cost by a factor of more than α, where $\alpha > 1$. Anshelevich et al. proved that there always exists a 3-approximate Nash equilibrium whose cost is equal to that of the optimum. Furthermore, they derived a polynomial time algorithm that gives a $(4.65 + \epsilon)$-approximate Nash equilibrium whose cost is twice the optimum.

In the following two paragraphs we describe the results known for network design games with Shapley cost sharing. The setting was introduced in a second paper by Anshelevich et al. [3] who first analyzed unweighted games. Using elegant potential function arguments based on a potential by Monderer and Shapley [20], the authors proved that every directed or undirected graph admits a Nash equilibrium and that the price of stability is upper bounded by $H(n)$. Here $H(n) = \sum_{i=1}^{n} 1/i$ is the nth Harmonic number, which is closely approximated by the natural logarithm, i.e. $\ln(n+1) \leq H(n) \leq \ln n + 1$. The upper bound of $H(n)$ on the price of stability is tight for directed graphs. For undirected graphs Anshelevich et al. [3] showed a lower bound of 4/3 on the price of stability; the lower bound construction uses two agents that have to establish a connection to a common destination. Additionally, Anshelevich et al [3] considered weighted games and showed the existence of Nash equilibria in two-agent games. For directed graphs they gave a lower bound of $\Omega(\max\{n, \log W\})$ on the price of stability, where W is the total weight of all the agents.

Chen and Roughgarden [7] further investigated weighted games in directed graphs. They showed that there are graphs that do not admit Nash equilibria. Chen and Roughgarden then demonstrated that, for any $\alpha = \Omega(\log w_{\max})$, α-approximate Nash equilibria do exist and that the price of stability is $O((\log W)/\alpha)$. Here $w_{\max}$ is the maximum weight of any agent. These trade-offs are nearly tight. Further work on unweighted games was presented by Fiat et al. [13] and Chekuri et al. [6].

All the above results hold for standard Nash equilibria without coordination. The concept of strong Nash equilibria has been the subject of extensive studies in the game theoretic literature. Recent research in game theory and economics has also investigated strong Nash equilibria in the context of networking problems. A survey article presenting literature on network formation in cooperative games was written by van den Nouweland [21]. More concretely, the existence of networks that are stable against changes in links by any coalition is examined in [17]. Furthermore, Andelman et al. [1] analyzed strong equilibria with respect to scheduling as well as a different class of network creation games in which links may be formed between any pair of agents. For the latter games, strong Nash equilibria achieve a constant price of anarchy.

We became aware that, independent of our work, very recently Epstein et al. [10] studied strong Nash equilibria for unweighted network design games in directed graphs. They assume that each agent has to connect a pair of terminals and consider Shapley as well as general cost sharing mechanisms. Epstein et al. observe that strong Nash equilibria do not always exist and then present topological characterizations for equilibrium existence. They show that if each agent has to connect a terminal to a common destination, each series parallel graph has a strong Nash equilibrium. If arbitrary terminal pairs are allowed, every extension parallel graph admits a strong Nash equilibrium when Shapley cost sharing is adopted. Furthermore Epstein et al. analyze the quality of strong Nash equilibria, showing a bound of $\Theta(\log n)$ on the price of anarchy for Shapley cost sharing and a bound of 1 for general cost sharing schemes when each agent has to connect to a common destination.

Our contribution: This paper presents an in-depth study of network design games with Shapley cost sharing when coordination among agents is allowed. We present upper and lower bounds on the price of anarchy achieved by strong Nash equilibria. We study scenarios with unrestricted coordination, i.e. coalitions of any size (or weight) may be formed; we also consider settings where the size (or weight) of a coalition is limited.

The first part of the paper addresses unweighted network design games. We first observe that there are graphs that do not admit strong Nash equilibria and then give a sufficient existence condition. More specifically, we show that α-approximate strong Nash equilibria exist in any directed or undirected graph, for any $\alpha \geq H(c)$, if coalitions of size up to c are allowed, $1 \leq c \leq n$. Again, $H(c)$ is the cth Harmonic number. An α-approximate strong Nash equilibrium, for $\alpha \geq 1$, is one where no coalition (of prescribed size or weight) can deviate such that every member of the coalition improves its cost by a factor of more than α.

We next prove that the price of anarchy of strong Nash equilibria is upper bounded by $H(n) \approx \ln n$, allowing coalitions of any size. This upper bound holds for any directed or undirected graph that admits a strong Nash equilibrium. Hence, using coordination, we achieve an exponential improvement in terms of the price of anarchy, compared to non-cooperative environments. We show that the upper bound of $H(n)$ is tight in directed graphs. For undirected graphs we develop a lower bound of $\Omega(\sqrt{\log n})$ on the price of anarchy. These results can be generalized to α-approximate strong Nash equilibria, for any $\alpha \geq 1$. In this case all the upper and lower bounds multiply by a factor of α. For the generalized setting that coalitions of size up to c are allowed, $1 \leq c \leq n$, we prove an upper bound of $\alpha \frac{n}{c} H(c)$ on the price of anarchy of α-approximate strong Nash equilibria. Again, this bound holds for any directed or undirected graph that admits an α-approximate strong Nash equilibrium, for some $\alpha \geq 1$, and not just for the range $\alpha \geq H(c)$. Suppose that $\alpha = 1$. If $c = 1$, we obtain the anarchy ratio of n achieved by standard equilibria. If $c = n$, we obtain the best ratio of $H(n)$. Since $H(n)$ is a lower bound on the price of stability of (standard) Nash equilibria in directed graphs [3], we conclude that in directed graphs the worst-case performance ratios of strong Nash equilibria are essentially always as good as the performance ratios achievable by the best standard Nash equilibria.

In the second part of the paper we extend the above results to weighted network design games. We first give a sufficient condition for the existence of α-approximate strong Nash equilibria. We then prove that in directed and undirected graphs the price of anarchy of strong Nash equilibria is at most $1 + \ln W$ if the formation of coalitions is not restricted. Here W is the sum of the weights of all agents. For directed graphs we show a matching lower bound of $\Omega(\log W)$. For undirected graphs we prove a lower bound of $\Omega(\sqrt{\log W})$. Again, for any $\alpha \geq 1$, the results extend to α-approximate strong Nash equilibria, where the lower and upper bounds simply multiply by α. When coordination among agents is limited, we consider two scenarios: (1) As usual, the number of agents in a coalition might be limited. (2) The sum of the weights of the agents forming a coalition may be limited so that agents of high weight cannot leave agents of low weight in costly configurations. For this general setting we present bounds trading the price of anarchy vs. the coalition size or weight. Furthermore, we prove a lower bound on the price of stability of standard Nash equilibria in undirected graphs. We construct a family of graphs in which the price of stability is $\Omega(\log W / \log \log W)$. No super-constant lower bound was known for undirected graphs, neither for weighted nor for unweighted games. Our lower bound holds even if every agent has to connect only a pair of terminals. However, individual terminal pairs are allowed. Together with the known lower bound of $\Omega(\log W)$ for directed graphs [3], we conclude that, in undirected as well as directed graphs, anarchy bounds of strong Nash equilibria essentially match the stability bounds of standard Nash equilibria.

We remark that our set of results is mostly disjoint from that by Epstein et al. [10]. The results provided in both [10] and this paper are the fact that strong Nash equilibria do not always exist, the upper bound of $\frac{n}{c}H(c)$ and the lower bound of $H(n)$ on the price of anarchy in unweighted games. While the upper bound proof by Epstein et al. is based on the potential function by Monderer and Shapley, in our paper we use new combinatorial arguments to establish the result. Generally speaking, our study here is more comprehensive in that we allow each agent to connect set of terminals, consider directed and undirected graphs as well as unweighted and weighted games.

Analysis techniques: As mentioned above, our upper bounds on the price of anarchy are achieved using new combinatorial arguments that do not rely on potential functions: Starting from a strong Nash equilibrium, we perform a sequence of specific strategy changes for varying size coalitions. For each strategy change there exists one unsatisfied agent whose original cost can be bounded relative to the optimum. From a technical point of view our strongest contribution are the lower bounds for undirected graphs. We present a new recursive framework for constructing lower bounds in network design games. Applying the recursive construction for varying parameters, we are able to obtain anarchy as well as stability bounds in both unweighted and weighted games. The protocol could also be applied to derive bounds for directed graphs but simpler constructions work in the directed case. While the same recursive framework can be applied to construct graphs for anarchy and stability bounds, the analyses of the graphs differ. To establish anarchy bounds we have to prove that no coalition can deviate, which turns out to be a non-trivial task because all possible coalitions and strategy changes over the recursive levels must be examined. To establish a stability bound, we have to show that no better Nash equilibria exist. In fact, we will prove that our graphs admit only one Nash equilibrium.

2 Problem statement and definitions

Network design games: Consider a graph $G = (V, E, c)$ with a non-negative cost function $c : E \mapsto \mathbb{R}_+^0$ defined on the edges. Graph G may be directed or undirected as we will study network design in both directed and undirected graphs. Associated with G are n selfish agents, each of which has certain connectivity requirements. More specifically, let $T_i \subseteq V$ be the set of terminals that agent i wishes to connect. If G is a directed graph, then for (selected) terminal pairs $t, t' \in T_i$ we additionally have to specify which direction between the pair should be established. A strategy of an agent i consists of a set $S_i \subseteq E$ of edges satisfying the connection requirements. If G is undirected, S_i is in fact a minimal tree connecting T_i. A combination $\mathcal{S}$ of strategies is the vector $\mathcal{S} = (S_1, \ldots, S_n)$ of individual agent strategies. Edges used by the agents have to be paid for. We consider Shapley cost sharing mechanisms that split the cost $c(e)$ of an edge e in a fair manner among the agents using that edge. In an *unweighted game*, if k agents use an edge e, then each of the k agents pays a share of $c(e)/k$ for that edge. Thus, for a combination $\mathcal{S}$ of strategies, the total cost of agent i is equal to $cost_i(\mathcal{S}) = \sum_{e \in S_i} c(e)/|\{j : e \in S_j\}|$. In a *weighted game* each agent i has a weight w_i and pays a share proportional to its weight. For any 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. Formally, the cost of agent i in a weighted game is $cost_i(\mathcal{S}) = \sum_{e \in S_i} c(e)w_i/W_e$.

Strong Nash equilibria: We are interested in stable solutions where agents have no incentive to deviate from their strategies. Previous work has considered Nash equilibria that are resilient to deviations of single agents. A weakness of Nash equilibria is their vulnerability to deviations by coalitions of agents. To overcome this problem, Aumann [4] defined the notion of *strong Nash equilibria*. A strong Nash equilibrium is resilient to deviations of coalitions, i.e. there exists no coalition of agents that can jointly change strategies such that every agent in the coalition has a strictly smaller cost. Formally, let I be a non-empty coalition of agents. For a combination $\mathcal{S}$ of strategies, let $\mathcal{S}_I$ be the projection of $\mathcal{S}$ on I, i.e. $\mathcal{S}_I$ are the strategies of agents $i \in I$. Similarly, $\mathcal{S}_{-I}$ represents the strategies of agents $i \notin I$. For coalition I, let $\mathcal{S}'_I$ be another choice of strategies. A combination $\mathcal{S}$ of strategies forms a strong Nash equilibrium if, for no non-empty coalition I, there exists a strategy change $\mathcal{S}'_I$ such that $cost_i(\mathcal{S}'_I, \mathcal{S}_{-I}) < cost_i(\mathcal{S})$, for all agents $i \in I$. Note that a standard Nash equilibrium is a strong Nash equilibrium where only coalitions of size one are allowed. In this spirit one can consider generalized settings in which coalitions of size at most c are permitted, $1 \leq c \leq n$. As for weighted games we will also be interested in scenarios where the total weight of agents forming a coalition is limited. This ensures that agents of high weight cannot impose too much control on agents outside a coalition.

As we shall see, strong Nash equilibria do not always exist. For this reason we relax the notion of stability, calling a combination of strategies stable if agents cannot improve their cost by a factor of more than α. More specifically, for a real value $\alpha \geq 1$, a combination $\mathcal{S}$ of strategies forms an α-approximate strong Nash equilibrium if, for no non-empty coalition I, there exists a strategy change $\mathcal{S}'_I$ such that $cost_i(\mathcal{S}'_I, \mathcal{S}_{-I}) < cost_i(\mathcal{S})/\alpha$, for all agents $i \in I$. Similarly, we can define α-approximate Nash equilibria when the size or weight of a coalition is limited. We remark that in the context of α-approximate strong equilibria another definition seems reasonable. We could call a combination of strategies an α-approximate strong Nash equilibrium if no coalition can improve its *total* cost by a factor of more than α, while still requiring that every agent of the coalition performs strictly better than before. Obviously, an α-approximate Nash equilibrium according this this second definition is an α-approximate equilibrium under the former definition, but not vice versa. Thus, our original definition allows for more configurations representing equilibrium states. For this reason and because our first definition requires a sufficiently high benefit for *each* agent of a coalition to perform a strategy change, we adopt the original definition in this paper. However, all the results that we will present in the following sections also hold for the second definition as well.

Performance measures: We are interested in the performance of strong Nash equilibria relative to the social optimum. Let $cost(\mathcal{S}) = \sum_{i=1}^{n} cost_i(\mathcal{S})$ be the total cost of all the agents and let $cost(OPT)$ be the cost of the globally optimal solution. We say that strong Nash equilibria *achieve a price of anarchy of* c if $\max_{\mathcal{S}} \frac{cost(\mathcal{S})}{cost(OPT)} \leq c$, where the maximum is taken over all strong Nash equilibria. The notion can be extended to (α-approximate) strong Nash equilibria with coalitions of limited size or weight. In this paper we will also be interested in the *price of stability* of standard Nash equilibria where coordination among agents is not allowed. The price of stability is $\min_{\mathcal{S}} \frac{cost(\mathcal{S})}{cost(OPT)}$, where the minimum is taken over all Nash equilibria.

3 Upper bounds for unweighted games

We study the existence of strong Nash equilibria and then develop upper bounds on the price of anarchy. The proof of the following proposition is presented in the Appendix.

Proposition 1 *There exist directed and undirected graphs that do not admit strong Nash equilibria.*

Theorem 1 *In any directed or undirected graph, α-approximate strong Nash equilibria exist, for any $\alpha \geq H(c)$, if coalitions of size up to c are allowed.*

Proof. We use a classical potential function by Monderer and Shapley [20] to show the existence of α-approximate strong Nash equilibria. Given a graph $G = (V, E, c)$ and a combination $\mathcal{S} = (S_1, \dots, S_n)$ of strategies, let n_e be the number of agents currently using edge $e \in E$ in their strategies, i.e. $n_e = |\{i : e \in S_i\}|$. The potential is defined as $\Phi(\mathcal{S}) = \sum_{e \in E} c(e) H(n_e)$. We will show that while $\mathcal{S}$ does not form an α-approximate strong Nash equilibrium, when allowing coalitions of size up to c, any α-*improvement move* strictly decreases the potential. An α-improvement move, for a coalition I with $|I| \leq c$, is a strategy change $\mathcal{S}'_I$ such that $cost_i(\mathcal{S}'_I, \mathcal{S}_{-I}) < cost_i(\mathcal{S})/\alpha$, for any agent $i \in I$. Suppose that we perform a sequence of such α-improvement moves starting from the social optimum. As the potential is upper bounded by $H(n) cost(OPT)$ and lower bounded by 0, the sequence of improvement moves must converge to an α-approximate strong Nash equilibrium.

We analyze an α-improvement move, performed by a coalition I with $|I| \leq c$. The strategy change $\mathcal{S}'_I$ of I can be viewed as being executed in two steps. (1) In a first step agents $i \in I$ drop all the edges used in strategies S_i. At this point no agent $i \in I$ shares the cost of any edge. (2) In a second step agents $i \in I$ join the edges they want to use in their new strategies $\mathcal{S}'_I$. Let E_1 be the set of edges dropped in step (1), and let E_2 be the set of edges added in step (2). These edge sets need not be disjoint. For any $e \in E$, let n_e^1 be the number of agents sharing e just after step (1) and let n_e^2 be the number of agents sharing e after step (2). The absolute value of the cost reduction experienced by I due to step (1) is

$$cost^- = \sum_{e \in E_1} c(e) \frac{n_e - n_e^1}{n_e},$$

because $e \in E_1$ is dropped by $n_e - n_e^1$ agents that each paid a share of $c(e)/n_e$. The value of this cost reduction is equal to the cost of I in the original configuration, i.e. $cost^- = \sum_{i \in I} cost_i(\mathcal{S})$, because after step (1) the cost of I is 0. The cost increase of I due to step (2) is

$$cost^+ = \sum_{e \in E_2} c(e) \frac{n_e^2 - n_e^1}{n_e^2},$$

because $e \in E_2$ is bought by $n_e^2 - n_e^1$ agents $i \in I$ who pay $c(e)/n_e^2$ each. This cost increase is equal to the cost of I in the new configuration, i.e. $cost^+ = \sum_{i \in I} cost_i(\mathcal{S}'_I, \mathcal{S}_{-I})$, because the cost of I was 0 before step (2) and the strategy change is complete after step (2). Using the definition of an α-improvement move we have $\sum_{i \in I} cost_i(\mathcal{S}'_I, \mathcal{S}_{-I}) < \sum_{i \in I} cost_i(\mathcal{S})/\alpha$ and hence

$$\alpha cost^+ - cost^- < 0. \tag{1}$$

Next we consider the potential change $\Delta\Phi$. The potential change stems from edges $e \in E_1 \cup E_2$ where cost shares change. Let Φ^- be the absolute value of the potential drop due to step (1) of the improvement move and let Φ^+ be the potential increase due to step (2). We will show $-\Phi^- \leq -cost^-$ and $\Phi^+ \leq \alpha cost^+$. This implies $\Delta\Phi = -\Phi^- + \Phi^+ \leq -cost^- + \alpha cost^+$ and using (1) we obtain $\Delta\Phi < 0$, which is to be proven.

To verify $-\Phi^- \leq -cost^-$ we observe

$$\Phi^- = \sum_{e \in E_1} c(e)(H(n_e) - H(n_e^1)) \geq \sum_{e \in E_1} c(e) \frac{n_e - n_e^1}{n_e} = cost^-.$$

The inequality holds because $H(n_e) - H(n_e^1) = 1/(n_e^1 + 1) + 1/(n_e^1 + 2) + \ldots + 1/n_e \geq (n_e - n_e^1)/n_e$. It remains to prove $\Phi^+ \leq \alpha cost^+$. The potential increase is given by $\Phi^+ = \sum_{e \in E_2} c(e)(H(n_e^2) - H(n_e^1))$. We show that for any $e \in E_2$,

$$H(n_e^2) - H(n_e^1) \leq H(n_e^2 - n_e^1) \frac{n_e^2 - n_e^1}{n_e^2}. \tag{2}$$

The desired inequality for the potential increase then follows because $n_e^2 - n_e^1 \leq c$ as at most c agents can join any edge in step (2) and $H(c) \leq \alpha$. The expressions in (2) are

$$\begin{aligned} H(n_e^2) - H(n_e^1) &= \frac{1}{n_e^1 + 1} + \frac{1}{n_e^1 + 2} + \ldots + \frac{1}{n_e^2} \\ H(n_e^2 - n_e^1) \frac{n_e^2 - n_e^1}{n_e^2} &= (1 + \frac{1}{2} + \ldots + \frac{1}{n_e^2 - n_e^1}) \frac{n_e^2 - n_e^1}{n_e^2}. \end{aligned}$$

We compare the kth terms of these expressions, for $k = 1, \ldots, n_e^2 - n_e^1$, and establish (2) by proving $\frac{1}{n_e^1 + k} \leq \frac{1}{k} \cdot \frac{n_e^2 - n_e^1}{n_e^2}$. This is equivalent to showing $0 \leq n_e^1(n_2^2 - n_e^1) - kn_e^1$, and this holds because $f(k) = n_e^1(n_2^2 - n_e^1) - kn_e^1$ is decreasing in k and $f(n_e^2 - n_e^1) = 0$. $\square$

Theorem 2 *In any directed or undirected graph and for any $\alpha \geq 1$, the price of anarchy of α-approximate strong Nash equilibria is upper bounded by $\frac{\alpha n}{c} H(c)$ if coalitions of size up to c are allowed.*

If there are no restrictions on the coalition size and we are interested in true strong Nash equilibria (i.e. $\alpha = 1$), we obtain:

Corollary 1 *In any directed or undirected graph the price of anarchy of strong Nash equilibria is upper bounded by $H(n)$.*

Proof of Theorem 2. Get G be a graph that admits α-approximate strong Nash equilibria, for some $\alpha \geq 1$, and let $\mathcal{S} = (S_1, \ldots, S_n)$ be such an equilibrium state. The basic idea of the proof is to consider all coalitions of size exactly c. For each coalition I we perform a process consisting of exactly c steps in which the agents of I try to buy the edges of the social optimum. At the end of each step exactly one agent will leave the process. Making use of the fact that in $\mathcal{S}$ no coalition of size up to c can improve its cost by a factor of more than α, we will be able to upper bound $cost_i(\mathcal{S})$ of the agent i leaving the process relative to the cost of the social optimum. More specifically, we will prove that for any coalition I of size exactly c,

$$\sum_{i \in I} cost_i(\mathcal{S}) \leq \alpha H(c) cost(OPT). \tag{3}$$

Let $\mathcal{I}$ be the set of all coalitions of size exactly c. Summing (3) over all the $\binom{n}{c}$ coalitions $I \in \mathcal{I}$, we obtain $\sum_{I \in \mathcal{I}} \sum_{i \in I} cost_i(\mathcal{S}) \leq \alpha \binom{n}{c} H(c) cost(OPT)$. Any fixed agent i, $1 \leq i \leq n$, occurs in exactly $\binom{n-1}{c-1}$ coalitions $I \in \mathcal{I}$. Hence $\sum_{I \in \mathcal{I}} \sum_{i \in I} cost_i(\mathcal{S}) = \binom{n-1}{c-1} cost(\mathcal{S})$. We conclude

$$cost(\mathcal{S}) \leq \binom{n}{c} / \binom{n-1}{c-1} \cdot \alpha H(c) cost(OPT) = \frac{n}{c} \cdot \alpha H(c) cost(OPT),$$

which establishes the stated price of anarchy.

Fix an arbitrary coalition I of size exactly c. We will prove (3). Let E^{OPT} be the set of edges bought by the social optimum and, for any $i \in I$, let E_i^{OPT} be a minimal set of edges necessary to connect the terminals of agent i within the optimal solution. We now start the process mentioned above. Let $I_1 := I$ be the initial coalition consisting of c agents. Suppose that we have already performed $k-1$ steps of the process, where initially $k = 1$, and let I_k be the coalition given at the beginning of the kth step, where $1 \leq k \leq c$. The kth step proceeds as follows. Starting from initial configuration $\mathcal{S}$, the agents of I_k perform a strategy change $\mathcal{S}_{I_k}^k$ in which $i \in I_k$ buys set E_i^{OPT}. Let $\mathcal{S}^k = (\mathcal{S}_{I_k}^k, \mathcal{S}_{-I_k})$ be the resulting configuration. The new cost of agent $i \in I_k$ is

$$cost_i(\mathcal{S}^k) = \sum_{e \in E_i^{OPT}} \frac{c(e)}{|\{j \in I_k : e \in E_j^{OPT}\} \cup \{j \notin I_k : e \in S_j\}|} \leq \sum_{e \in E_i^{OPT}} \frac{c(e)}{|\{j \in I_k : e \in E_j^{OPT}\}|}.$$

Since the original configuration $\mathcal{S}$ forms an α-approximate strong Nash equilibrium, the strategy change cannot improve the cost of every agent $i \in I_k$ by a factor of more than α. Thus there must exist an agent i_k with $cost_{i_k}(\mathcal{S}^k) \geq cost_{i_k}(\mathcal{S})/\alpha$ and hence

$$cost_{i_k}(\mathcal{S}) \leq \alpha \sum_{e \in E_{i_k}^{OPT}} \frac{c(e)}{|\{j \in I_k : e \in E_j^{OPT}\}|}. \tag{4}$$

This agent i_k leaves the coalition I_k. If there is more than one agent satisfying the above cost inequality, we select an arbitrary of them. The new coalition at the end of the step is $I_{k+1} := I_k \setminus \{i_k\}$. The process ends after exactly c steps when the coalition is empty. Summing (4) over all the c steps, taking into account that the sequence of agents leaving the process forms I, we find

$$\sum_{i \in I} cost_i(\mathcal{S}) \leq \alpha \sum_{k=1}^{c} \sum_{e \in E_{i_k}^{OPT}} \frac{c(e)}{|\{j \in I_k : e \in E_j^{OPT}\}|}. \tag{5}$$

We analyze the right-hand side of the above inequality, which sums edge costs $c(e)$ over edges $e \in E^{OPT}$. Consider any fixed edge $e \in E^{OPT}$ and let $n_e = |\{i \in I : e \in E_i^{OPT}\}|$ be the number of agents in I using e in the described strategy changes. The cost of e contributes to the right-hand side of (5) whenever one of the n_e agents leaves the process. The ℓth time this happens, the contribution is $c(e)/(n_e - \ell + 1)$, for $\ell = 1, \ldots, n_e$. Thus, the cost contribution is $c(e)H(n_e) \leq c(e)H(c)$ and we conclude $\sum_{i \in I} cost_i(\mathcal{S}) \leq \alpha \sum_{e \in E^{OPT}} c(e)H(c) = \alpha H(c) cost(OPT)$. □

4 Lower bounds for unweighted games

We first present a lower bound for directed graphs. This lower bound implies that if there is no restriction on the coalition size, our upper bound of Corollary 1 is optimal.

Theorem 3 *In directed graphs and for any $\alpha \geq 1$, the price of anarchy of α-approximate strong Nash equilibria is at least $\alpha \max\{n/c, H(n)\}$ if coalitions of size at most c are allowed.*

Proof. We modify lower bound graphs that were presented previously in the literature [2]. For the bound of $\alpha n/c$, consider a simple graph consisting of two vertices s and t that are connected by two parallel edges of cost αn and $c + \epsilon$, respectively, see Figure 1(a). Associated with the graph are n agents, all of which have to connect terminals s and t. An optimal solution will buy the edge of cost $c + \epsilon$. On the other hand, the configuration in which all the n agents share the expensive edge of cost αn, each one paying a cost of α, represents an α-approximate strong Nash equilibrium: Any coalition of size up to c, when performing a strategy change and buying the edge of cost $c + \epsilon$, incurs a cost of at least $1 + \epsilon/c$ per agent. Hence the agents of the coalition do not save a factor of more than α in cost.

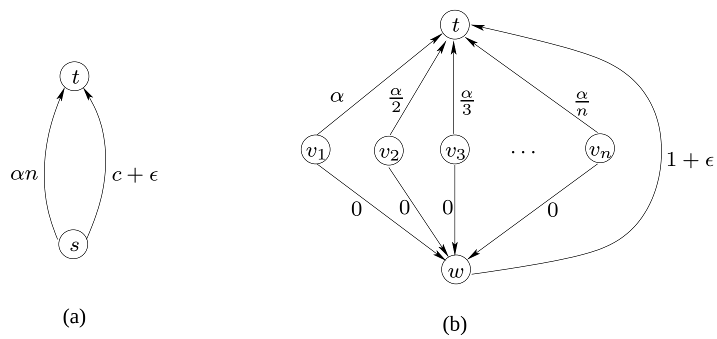

Figure 1: Directed graphs enforcing a high price of anarchy.

In order to establish the lower bound of $\alpha H(n)$, we use the graph depicted in Figure 1(b). There are n vertices $v_1, \ldots, v_n$, where v_i is connected to a vertex t via a directed edge (v_i, t) of cost α/i and to a vertex w via a directed edge (v_i, w) of cost 0. Additionally, there is a directed edge (w, t) of cost $1 + \epsilon$. Associated with the graph are n agents, where agent i has to connect v_i to t. An optimal solution satisfies the connection requirements by buying the edges of cost 0 and the edge (w, t) of cost $1 + \epsilon$. The configuration in which agent i connects v_i to t using its private edge (v_i, t) of cost α/i forms an α-approximate strong Nash equilibrium. Any coalition of size, say c, that performs a strategy change and purchases edge (w, t) incurs a cost of $(1+\epsilon)/c$ per agent. However, there is at least one agent in the coalition whose original cost was at most α/c and to whom the incentive of changing is not sufficiently high. $\square$

We next develop a lower bound for undirected networks. Our lower bound construction is quite involved and we therefore concentrate on the most general scenario where there is no limit on the coalition size.

Theorem 4 *For any $\alpha \geq 1$, there exists a family of undirected graphs, each admitting an α-approximate strong Nash equilibrium whose cost is $\Omega(\alpha\sqrt{\log n})$ times that of the social optimum.*

Proof. For ease of exposition we first prove the theorem for $\alpha = 1$ and then show how to adapt the proof for any $\alpha > 1$. We present a recursive definition of graphs G. Let n be a positive integer such that $\lfloor\sqrt{\log n}\rfloor \geq 2$. In this proof logarithms are taken to the base 3. Let $d_{\max} = \lfloor\sqrt{\log n}\rfloor - 1$. The recursive construction proceeds in $d_{\max}+1$ steps. At the bottom level of the recursion, i.e. at maximum depth $d_{\max}$, G consists of graphs $G^{d_{\max}}$

of order $d_{\max}$, cf. Figure 2(a). A graph $G^{d_{\max}}$ is composed of a *stem edge* $\{v, w\}$ of cost $s_{d_{\max}} = 1/3^{d_{\max}}$ and a *bridge* $\{u, v\}$ *of order* $d_{\max}$ having cost $b_{d_{\max}} = 2/3^{2d_{\max}}$. The bridge and the stem are joined at vertex v. Vertices u and w are connected via an *arc* $\{u, w\}$ *of order* $d_{\max}$ having cost $a_{d_{\max}} = 1/3^{d_{\max}}$. We call u the *base* and w the *tip* of $G^{d_{\max}}$. Associated with $G^{d_{\max}}$ are $n_{d_{\max}} = \lceil n/3^{d_{\max}(d_{\max}+1)} \rceil$ *agents of order* $d_{\max}$, each having to connect terminals u and w. By the choice of $d_{\max}$ we have $n_{d_{\max}} \geq 1$.

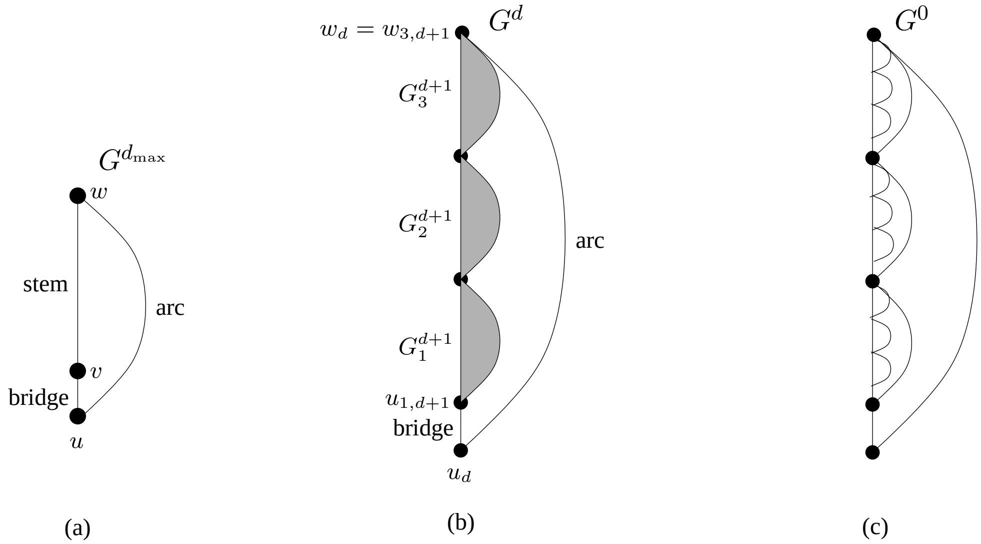

Figure 2: The recursive construction of graphs G^d.

Assume that graphs of order $d_{\max}, d_{\max} - 1, \ldots, d+1$ are defined. Then a graph G^d *of order* d, which resides a depth d of the recursion, is constructed as follows, see Figure 2(b). Graph G^d consists of three graphs G_1^{d+1}, G_2^{d+1} and G_3^{d+1} of order $d+1$ that are attached to each other. More specifically, the tip of G_1^{d+1} and the base of G_2^{d+1} are merged, i.e. the two vertices are united, and the tip of G_2^{d+1} is merged with the base of G_3^{d+1}. Let $u_{1,d+1}$ be the base of G_1^{d+1}. Attached to this vertex is a bridge $\{u_d, u_{1,d+1}\}$ of order d having a cost of $b_d = 2/3^{2d}$. Let $w_{3,d+1}$ be the tip of G_3^{d+1} and set $w_d := w_{3,d+1}$. We call u_d the *base* and w_d the tip of G^d. Additionally, G^d contains an arc $\{u_d, w_d\}$ of order d connecting the base and the tip. This arc has cost $a_d = 1/3^d$. Associated with G^d are $n_d = \lceil n/3^{d(d+1)} \rceil - 3\lceil n/3^{(d+1)(d+2)} \rceil$ agents of order d, all of which have to connect u_d to w_d. As we shall see, these agents will govern the connection decisions within G^d. The bridge will have the effect that in a strong Nash equilibrium, the order-d agents will establish their connections using the arc of order d instead of routing through the graphs G_k^{d+1}, $1 \leq k \leq 3$.

The construction proceeds down to a depth $d = 0$. Associated with graph G^0 are $n_0 = \lceil n/3^0 \rceil - 3\lceil n/3^2 \rceil = n - 3\lceil n/3^2 \rceil$ agents of order 0 that have to connect the outermost vertices of G^0. Graph $G := G^0$ is the graph we will work with. A high level sketch of $G = G^0$ is given in Figure 2(c).

We start with some observations on $G = G^0$. First, all the vertices and terminals of the graph are located on a *backbone* consisting of all the stem edges and bridges. The nested structure of G^0 contains 3^d subgraphs of order d, for any $0 \leq d \leq d_{\max}$.

Proposition 2 *The least expensive path connecting the base and the tip of a graph* G^d *using only edges of* G^d *has a total edge cost of exactly* $1/3^d$*, for any* $0 \leq d \leq d_{\max}$.

Proof. The statement holds for $d = d_{\max}$ as the arc of $G^{d_{\max}}$ has cost $a_{d_{\max}} = 1/3^{d_{\max}}$ while the path crossing the bridge has cost $b_{d_{\max}} + s_{d_{\max}} = 2/3^{2d_{\max}} + 1/3^{d_{\max}}$. Suppose that the statement holds for depths $d_{\max}, \ldots, d+1$. In G^d the arc of order d has cost $a_d = 1/3^d$ while, using the inductive hypothesis, any path using the bridge of order d has a cost of at least $b_d + 3 \cdot 1/3^{d+1} = 2/3^{2d} + 1/3^d > 1/3^d$. □

The total number of agents associated with G^0 and all of its subgraphs is equal to

$$
\begin{aligned}
N_0 &= \sum_{d=0}^{d_{\max}-1} 3^d(\lceil n/3^{d(d+1)}\rceil - 3\lceil n/3^{(d+1)(d+2)}\rceil) + 3^{d_{\max}}\lceil n/3^{d_{\max}(d_{\max}+1)}\rceil \\
&= n - 3^{d_{\max}}\lceil n/3^{d_{\max}(d_{\max}+1)}\rceil + 3^{d_{\max}}\lceil n/3^{d_{\max}(d_{\max}+1)}\rceil \\
&= n.
\end{aligned}
$$

More generally, in $G = G^0$ the total number of agents associated with all the order-d graphs G^d and the subgraphs therein is, for any d with $0 \leq d \leq d_{\max}$,

$$
N_d = \sum_{i=d}^{d_{\max}-1} 3^i(\lceil n/3^{i(i+1)}\rceil - 3\lceil n/3^{(i+1)(i+2)}\rceil) + 3^{d_{\max}}\lceil n/3^{d_{\max}(d_{\max}+1)}\rceil = 3^d\lceil n/3^{d(d+1)}\rceil,
$$

which is equal to $n/3^{d^2}$ when ignoring ceilings.

The social optimum in G buys the backbone of the graph. As there are $3^{d_{\max}}$ graphs of order $d_{\max}$, the total cost of the stem edges is $3^{d_{\max}} s_{d_{\max}} = 3^{d_{\max}} \cdot 1/3^{d_{\max}} = 1$. There are 3^d graphs of order d, $0 \leq d \leq d_{\max}$, and hence the total cost of order-d bridges is $3^d b_d = 3^d \cdot 2/3^{2d} = 2/3^d$. Summing over all d we find that the total cost of the bridges is $\sum_{d=0}^{d_{\max}} 2/3^d \leq 3$. We conclude that the cost of the social optimum is bounded by 4.

Consider the configuration $\mathcal{S}$ in which, for any graph G^d within G, any order-d agent associated with this graph G^d establishes its required connection via the corresponding arc of order d. That is, $\mathcal{S}$ buys all the arcs. As we will show in the remainder of this proof, $\mathcal{S}$ forms a strong Nash equilibrium. We evaluate the cost of $\mathcal{S}$. As there are 3^d graphs of order d, the total cost of order-d arcs is $3^d a_d = 3^d \cdot 1/3^d = 1$, for any fixed d with $0 \leq d \leq d_{\max}$. Summing over all d, we obtain $cost(\mathcal{S}) = d_{\max} + 1 \geq \lfloor\sqrt{\log n}\rfloor$, and this establishes the desired performance ratio.

It remains to show that $\mathcal{S}$ is indeed a strong Nash equilibrium. To this end we have to show that no coalition I of agents has an *improvement move*. We will always consider non-empty coalitions. An improvement move, for a coalition I, is a strategy change $\mathcal{S}'_I$ such that $cost_i(\mathcal{S}'_I, \mathcal{S}_{-I}) < cost_i(\mathcal{S})$, for any agent $i \in I$. In our graph G, as all the agents have to connect pairs of terminals, a strategy of an agent is a simple path connecting the desired vertices. The property that there exists no improvement move follows from Lemma 1, which we prove in the sequel.

Lemma 1 *For $d = 0, \ldots, d_{\max}$, no coalition involving agents of order d has an improvement move.*

For the proof of Lemma 1 we need Lemma 2 which we prove first.

Lemma 2 *Consider a fixed d, $0 \leq d \leq d_{\max}$, and suppose that no coalition involving agents of order smaller than d has an improvement move. Furthermore, assume that no coalition I involving agents of order d has an improvement move in which an order-d agent $i \in I$ associated with a graph $G^d(i)$ chooses a path containing edges outside $G^d(i)$. Then no coalition involving agents of order d has an improvement move.*

Proof of Lemma 2. Let I be a coalition that involves agents of order d. We have to show that I has no improvement move. Based on the assumptions of the lemma, we can restrict ourselves to coalitions I that do not contain agents of order smaller than d. Furthermore, based on the assumptions, we only have to consider strategy changes where each order-d agent $i \in I$ establishes the required connection within its graph $G^d(i)$. Let $I' \subseteq I$ be any maximal sub-coalition of order-d agents that are associated with the same graph $G^d(I')$. We will show that any strategy change $\mathcal{S}'_{I'}$ that consists in choosing connection paths within $G^d(I')$ leads to a strictly higher cost for that sub-coalition, i.e. at least one agent $i \in I'$ has a strictly higher cost and the strategy change is no improvement move.

Graph $G^d(I')$ has n_d order-d agents associated with it. Let f be the fraction defecting, i.e. $f = |I'|/n_d$. In the original configuration $\mathcal{S}$, when routing through the arc of order d, sub-coalition I' paid a cost of $fa_d = f/3^d$. When changing strategy and choosing a different connection route within $G^d(I')$, each $i \in I'$ selects a path P_i that crosses the bridge of order d and then, if $d = d_{\max}$, traverses the stem edge of $G^{d_{\max}}(I')$ (see Fig. 2(a)). If $d < d_{\max}$, path P_i then traverses the order-$(d+1)$ graphs $G_k^{d+1}(I')$, $1 \le k \le 3$, located within $G^d(I')$ (see Fig. 2(b)). If $d = d_{\max}$, then the total cost of edges on P_i is $b_{d_{\max}} + s_{d_{\max}} \ge (1 + 2/3^d)/3^d$. If $d < d_{\max}$, then the total cost is at least $b_d + 3 \cdot 1/3^{d+1} \ge (1 + 2/3^d)/3^d$ because, by Proposition 2, the least expensive path traversing an order-$(d+1)$ graph has cost $1/3^{d+1}$. In both cases we have the same lower bound on the cost, expressed in terms of d. The cost of P_i is not shared by agents of order smaller than d, as they are not part of the original coalition I, nor is the cost shared by order-d agents associated with other graphs $G^d \neq G^d(I')$. The cost of P_i can only be shared by agents of order larger than d, and there exist N_{d+1} such agents if $d < d_{\max}$. If $d = d_{\max}$, the cost is not shared by other agents.

If $d = d_{\max}$, we are done because the new cost of I' is $(1 + 2/3^d)/3^d$, while the original cost was $f/3^d \le 1/3^d$. If $d < d_{\max}$, then at best all the N_{d+1} agents of order larger than d support the edges traversed by I' and the new cost of I' is at least $cost'_{I'} \ge \frac{fn_d}{fn_d + N_{d+1}}(1 + \frac{2}{3^d})\frac{1}{3^d}$. We will to show that $cost'_{I'}$ is higher than the original cost of $f/3^d$, which is equivalent to proving $\frac{n_d}{fn_d + N_{d+1}}(1 + \frac{2}{3^d}) > 1$. Since $0 < f \le 1$ it suffices to show

$$\frac{n_d}{n_d + N_{d+1}}(1 + \frac{2}{3^d}) > 1.$$

Using the definition of n_d, eliminating ceilings, we find

$$\begin{aligned} n_d &> n/3^{d(d+1)} - n/3^{(d+1)^2+d} - 3 = \frac{n3^d}{3^{(d+1)^2}} \left(3 - \frac{1}{3^{2d}} - \frac{3^{d^2+d+2}}{n}\right) \\ &\ge \frac{n3^d}{3^{(d+1)^2}} \left(2 - \frac{3^{(d+1)(d+2)}}{n}\right) \\ &\ge \frac{n3^d}{3^{(d+1)^2}}. \end{aligned}$$

The second inequality holds because $1/3^{2d} \le 1$. For the third inequality note that $d < d_{\max}$ and $d_{\max} = \lfloor\sqrt{\log n}\rfloor - 1$ imply $(d+1)(d+2) \le \log n$ and hence $3^{(d+1)(d+2)} \le n$. Moreover, we have $N_{d+1} < 2n/3^{(d+1)^2}$. We conclude

$$\frac{n_d}{n_d + N_{d+1}}(1 + \frac{2}{3^d}) > \frac{3^d}{3^d + 2}(1 + \frac{2}{3^d}) = \frac{3^d}{3^d + 2} \cdot \frac{3^d + 2}{3^d} = 1.$$

□

Proof of Lemma 1. We prove the lemma inductively for increasing values of d. For $d = 0$, the statement follows immediately from Lemma 2 as the assumptions of that lemma are trivially satisfied: There are no agents of order smaller than 0 and an agent of order 0 cannot connect its terminals using edges outside G^0. Suppose that the statement of Lemma 1 holds for depth $0, \ldots, d-1$. We prove that no coalition I involving order-d agents has an improvement move in which an order-d agent $i \in I$ associated with a graph $G^d(i)$ chooses a path using edges outside $G^d(i)$. The inductive step then follows from Lemma 2.

So consider a coalition I involving agents of order d and a corresponding strategy change $\mathcal{S}'_I$ in which at least one order-d agent chooses edges outside its order-d graph to connect the desired terminals. We show that the strategy change is not an improvement move. By the inductive hypothesis, we can restrict ourselves to coalitions I not involving agents of order smaller than d. Thus I only contains agents of order d or larger. Let $I' \subseteq I$ be the maximum sub-coalition of order-d agents i choosing connection paths outside their graph

$G^d(i)$. As $d \geq 1$, each such graph belongs to a graph G^{d-1} in the nested structure of G^0. Consider all graphs of order $d-1$ containing at least one $G^d(i)$, $i \in I'$, and number these order-$(d-1)$ graphs in an arbitrary way. Let J be the resulting index set. Each graph $G^{d-1,j}$, $j \in J$, contains three graphs $G_1^{d,j}$, $G_2^{d,j}$, $G_3^{d,j}$ of order d. For $k = 1, 2, 3$, let f_k^j be the fraction of the order-d agents associated with $G_k^{d,j}$ that are member of I', i.e. $f_k^j = |\{i \in I' : i \text{ is order-}d \text{ agent associated with } G_k^{d,j}\}|/n_d$. Recall that n_d is the number of agents associated with an order-d graph. We have $0 \leq f_k^j \leq 1$ and $f_1^j + f_2^j + f_3^j > 0$.

In the original configuration $\mathcal{S}$, coalition I' pays a total cost of $cost_{I'} = \sum_{j \in J} \sum_{k=1}^{3} f_k^j a_d$ because the order-d arcs are bought. This expression holds even if some of the f_k^j are zero. We show in the following that the new cost $cost'_{I'}$ of I' is strictly higher than $cost_{I'}$. Hence at least one agent in I' does not improve its cost and the strategy change is no improvement move.

In order to estimate $cost'_{I'}$ consider an agent $i \in I'$ and let $G^{d-1,j}$, with $j \in J$, be the graph where i's graph $G^d(i)$ is located. First suppose that $G^d(i) = G_1^{d,j}$, cf. Figure 3(a). After the strategy change, i connects the base and the tip of $G_1^{d,j}$ on a path P_i that uses edges outside $G_1^{d,j}$. Since strategies are simple paths, all the edges of P_i are outside $G_1^{d,j}$. Starting at the base of $G_1^{d,j}$, path P_i has to traverse the bridge of order $d-1$ in $G^{d-1,j}$. To reach the tip of $G_1^{d,j}$, path P_i has to travel to the tip of $G^{d-1,j}$. This can be done using the arc of order $d-1$ or using another subpath outside $G^{d-1,j}$. After having reached the tip of $G^{d-1,j}$, path P_i traverses $G_3^{d,j}$ and $G_2^{d,j}$, reaching the desired terminal. Ignoring edges visited between the base and the tip of $G^{d-1,j}$, agent i has to pay a share for the order-$(d-1)$ bridge and for the subpaths of P_i within $G_2^{d,j}$ and $G_3^{d,j}$. The total cost of edges traversed within graph $G_k^{d,j}$, $k \in \{2, 3\}$, is at least $1/3^d$ by Proposition 2. The cost may be shared with other agents.

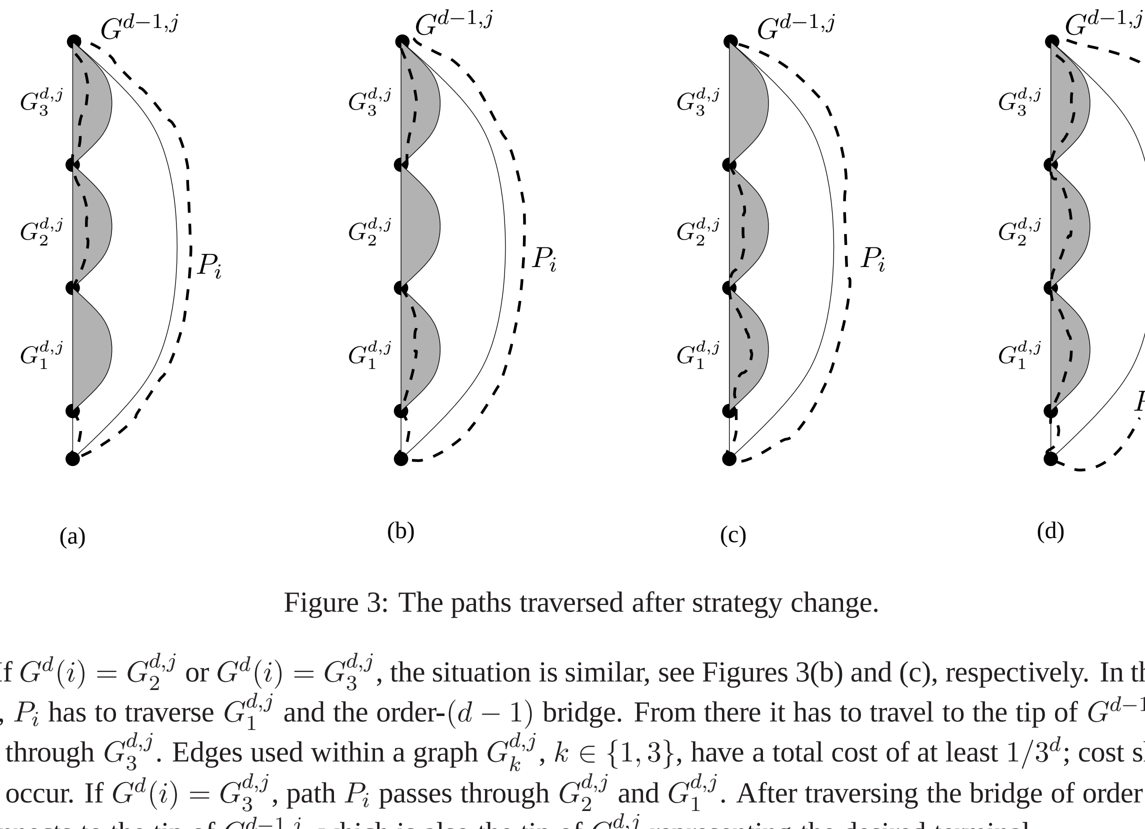

Figure 3: The paths traversed after strategy change.

If $G^d(i) = G_2^{d,j}$ or $G^d(i) = G_3^{d,j}$, the situation is similar, see Figures 3(b) and (c), respectively. In the first case, P_i has to traverse $G_1^{d,j}$ and the order-$(d-1)$ bridge. From there it has to travel to the tip of $G^{d-1,j}$ and pass through $G_3^{d,j}$. Edges used within a graph $G_k^{d,j}$, $k \in \{1, 3\}$, have a total cost of at least $1/3^d$; cost sharing may occur. If $G^d(i) = G_3^{d,j}$, path P_i passes through $G_2^{d,j}$ and $G_1^{d,j}$. After traversing the bridge of order $d-1$ it connects to the tip of $G^{d-1,j}$, which is also the tip of $G_3^{d,j}$ representing the desired terminal.

Next let C^j be the total number of agents $i' \in I'$ that are *not* associated with $G_1^{d,j}$, $G_2^{d,j}$ or $G_3^{d,j}$ but choose edges of these graphs when performing the strategy change. In order to use such edges, the new path $P_{i'}$ of i'

must pass through the base and the tip of $G^{d-1,j}$, see Figure 3(d). The path between these two vertices crosses the bridge of order $d-1$ and must consist of subpaths within $G_k^{d,j}$, $k = 1, 2, 3$.

We are ready to lower bound the new cost $cost'_{I'}$ of I'. To this end we will only consider the cost spent in graphs $G_k^{d,j}$, $1 \leq k \leq 3$ and $j \in J$, and on order-$(d-1)$ bridges in $G^{d-1,j}$. Fix a $j \in J$. Graph $G_1^{d,j}$ is traversed by exactly $(f_2^j + f_3^j)n_d + C^j$ agents from I', each using edges of cost at least $1/3^d$. The cost is shared by at most $(f_2^j + f_3^j)n_d + C^j + (1 - f_1^j)n_d + N_{d+1}$ agents. Here $(1 - f_1^j)n_d$ is the number of order-d agents associated with $G_1^{d,j}$ that establish their connection within this graph and N_{d+1} is the total number of agents of order larger than d that may participate in the strategy change and reside in the entire coalition I. If $d = d_{\max}$, the we set $N_{d+1} = 0$. Thus sub-coalition I' spends a cost of at least

$$\frac{1}{3^d} \cdot \frac{(f_2^j + f_3^j)n_d + C^j}{(1 - f_1^j + f_2^j + f_3^j)n_d + N_{d+1} + C^j}$$

in graph $G_1^{d,j}$. Similarly, in $G_2^{d,j}$ and $G_3^{d,j}$ the costs are

$$\frac{1}{3^d} \cdot \frac{(f_1^j + f_3^j)n_d + C^j}{(1 - f_2^j + f_1^j + f_3^j)n_d + N_{d+1} + C^j} \quad \text{and} \quad \frac{1}{3^d} \cdot \frac{(f_1^j + f_2^j)n_d + C^j}{(1 - f_3^j + f_1^j + f_2^j)n_d + N_{d+1} + C^j}.$$

Finally the bridge of order $d-1$ has cost $b_{d-1} = 2/3^{2(d-1)}$, which is shared by $(f_1^j + f_2^j + f_3^j)n_d + N_{d+1} + C^j$ agents and I' incurs a cost of $\frac{2}{3^{2(d-1)}}((f_1^j + f_2^j + f_3^j)n_d + C^j)/((f_1^j + f_2^j + f_3^j)n_d + N_{d+1} + C^j)$.

Note that for any $k \in \{1, 2, 3\}$ the other two indices from that set can be expressed as $k' = k \bmod 3 + 1$ and $k'' = (k+1) \bmod 3 + 1$. We conclude

$$\begin{aligned} cost'_{I'} \quad \geq \quad & \sum_{j \in J} \left(\sum_{k=1}^{3} \frac{(f_{k'}^j + f_{k''}^j)n_d + C^j}{(1 - f_k^j + f_{k'}^j + f_{k''}^j)n_d + N_{d+1} + C^j} \cdot \frac{1}{3^d} \right. \\ & \left. + \frac{(f_1^j + f_2^j + f_3^j)n_d + C^j}{(f_1^j + f_2^j + f_3^j)n_d + N_{d+1} + C^j} \cdot \frac{2}{3^{2(d-1)}} \right). \end{aligned}$$

Ratios of the form $(x+c)/(y+c)$ are increasing in c if $x \leq y$. Hence we can drop the terms C^j and obtain

$$cost'_{I'} \geq \sum_{j \in J} \left(\sum_{k=1}^{3} \frac{(f_{k'}^j + f_{k''}^j)n_d}{(1 - f_k^j + f_{k'}^j + f_{k''}^j)n_d + N_{d+1}} \cdot \frac{1}{3^d} + \frac{(f_1^j + f_2^j + f_3^j)n_d}{(f_1^j + f_2^j + f_3^j)n_d + N_{d+1}} \cdot \frac{2}{3^{2(d-1)}} \right),$$

Reordering the expression in the brackets, by focusing on one particular f_k^j in the numerators, we find

$$\begin{aligned} cost'_{I'} \quad \geq \quad & \sum_{j \in J} \sum_{k=1}^{3} \left(\left(\frac{f_k^j n_d}{(1 + f_k^j - f_{k'}^j + f_{k''}^j)n_d + N_{d+1}} + \frac{f_k^j n_d}{(1 + f_k^j + f_{k'}^j - f_{k''}^j)n_d + N_{d+1}} \right) \cdot \frac{1}{3^d} \right. \\ & \left. + \frac{f_k^j n_d}{(f_1^j + f_2^j + f_3^j)n_d + N_{d+1}} \cdot \frac{2}{3^{2(d-1)}} \right). \end{aligned}$$

To simplify the last expression we observe that for any real values x, y and c inequality $\frac{1}{x-y+c} + \frac{1}{x+y+c} \geq \frac{2}{x+c}$ holds, which we apply for $x = f_k^j n_d$ and $y = (f_{k'}^j - f_{k''}^j)n_d$ as well as $c = n_d + N_{d+1}$. Furthermore,

$$\frac{f_k^j n_d}{(f_1^j + f_2^j + f_3^j)n_d + N_{d+1}} \geq \frac{f_k^j n_d}{(f_k^j + 2)n_d + N_{d+1}} \geq \frac{1}{2} \cdot \frac{f_k^j n_d}{(1 + f_k^j)(n_d + N_{d+1})}.$$

Hence

$$\begin{aligned} cost'_{I'} &\geq \sum_{j\in J}\left(\sum_{k=1}^{3}\frac{2f_k^j n_d}{(1+f_k^j)n_d+N_{d+1}}\cdot\frac{1}{3^d}+\frac{f_k^j n_d}{(1+f_k^j)(n_d+N_{d+1})}\cdot\frac{1}{3^{2(d-1)}}\right) \\ &> \sum_{j\in J}\left(\sum_{k=1}^{3}\frac{2f_k^j n_d}{(1+f_k^j)(n_d+N_{d+1})}(1+\frac{2}{3^d})\frac{1}{3^d}\right). \end{aligned}$$

As shown at the end of the proof of Lemma 2, $\frac{n_d}{n_d+N_{d+1}}(1+\frac{2}{3^d})>1$ if $d<d_{\max}$. If $d=d_{\max}$, then $N_{d+1}=0$ and the inequality is also satisfied. In each case

$$cost'_{I'}>\sum_{j\in J}\sum_{k=1}^{3}\frac{2f_k^j}{1+f_k^j}\frac{1}{3^d}\geq\sum_{j\in J}\sum_{k=1}^{3}f_k^j\frac{1}{3^d}$$

and the new cost of I' is strictly larger than the original cost of I' in configuration $\mathcal{S}$. □

This completes the proof of Theorem 4 for $\alpha=1$. We finally show how to adapt the proof for any $\alpha>1$. In the construction of the graphs G^d only the costs of the arcs change. An arc of order d now has cost αa_d. All other costs remain the same. This increases the cost of configuration $\mathcal{S}$ by a factor of α, i.e. $cost(\mathcal{S})\geq\alpha\lfloor\log n\rfloor$ while the cost of the social optimum remains the same. This establishes a performance ratio of $\Omega(\alpha\sqrt{\log n})$.

In the statements of Lemmas 1 and 2, the term “improvement move” has to be replaced by “α-improvement move”. An α-improvement move, for a coalition I, is a strategy change $\mathcal{S}'_I$ such that $cost_i(\mathcal{S}'_I,\mathcal{S}_{-I})<cost_i(\mathcal{S})/\alpha$, for any agent $i\in I$. In the proof of Lemma 2 we considered any coalition I involving agents of order d or larger and investigated strategy changes where order-d agents establish connections with their respective graph of order d. We identified a sub-coalition I', with $f=|I'|/n_d$, incurring a new cost of $cost'_{I'}>f/3^d$. This cost inequality still hold in our modified graph as edge costs did not decrease. Since $cost'_{I'}>f/3^d=(\alpha f/3^d)/\alpha$ and $\alpha f/3^d$ is the original cost of I' in the scaled graph, the strategy change is no α-improvement move.

In the proof of Lemma 2 we studied coalitions I involving agents of order d or larger. We analyzed strategy changes in which order-d agents buy edges outside their graph of order d and identified a sub-coalition I' of order-d agents incurring a new total cost of $cost'_{I'}>\sum_{j\in J}\sum_{k=1}^{3}f_k^j/3^d$, where $G_k^{d,j}$ were the graphs the agents $i\in I'$ are associated with. Again, when arcs are scaled by a factor of α, this cost inequality still holds. As the original cost of I' in the scaled graph is $\sum_{j\in J}\sum_{k=1}^{3}\alpha f_k^j/3^d$, the strategy change is no α-improvement move. □

5 Weighted games

In this section we study weighted network design games where each agent i has a positive weight w_i. We scale the weights such that the minimum weight is equal to 1 and hence $w_i\geq 1$ for all agents. Let $W=\sum_{i=1}^{n}w_i$ be the total weight of all the agents.

If agents are allowed to coordinate their strategies, two scenarios are of interest. In a first setting we assume that coalitions of size up to c are allowed, for any $1\leq c\leq n$. In this case let W^c be the maximum total weight of any coalition having size at most c. In a second setting we assume that the total weight of a coalition is upper bounded so that agents of high weight cannot impose too much control on agents of low weight. In this case let $W^c_{\max}$ be the maximum total weight any coalition may have.

We extend our results shown for unweighted games.

5.1 Upper bounds

We first give a sufficient condition for the existence of strong Nash equilibria in weighted games and evaluate their performance in terms of the price of anarchy.

Theorem 5 *In any directed or undirected graph α-approximate strong Nash equilibria exist, for any $\alpha \geq 1+\ln(1+\overline{W})$. Here $\overline{W} = W^c$ if coalitions of size up to c are allowed and $\overline{W} = W^c_{\max}$ if coalitions of weight up to $W^c_{\max}$ are allowed.*

Proof. We use again potential function arguments to show the existence of α-approximate strong Nash equilibria but have to work with a more general potential function, compared to that used in unweighted games. Given a graph $G = (V, E, c)$ and a configuration $\mathcal{S} = (S_1, \ldots, S_n)$, let $E_{\mathcal{S}} = \cup_{i=1}^n S_i$ be the union of all edges used by the agents. For any $e \in E_{\mathcal{S}}$, let $W_e = \sum_{i:e\in S_i} w_i$ be the total weight of the agents currently using e in their strategies. Define

$$\Phi(\mathcal{S}) = \sum_{e\in E_{\mathcal{S}}} c(e)(1 + \ln W_e).$$

We show that while $\mathcal{S}$ does not form an α-approximate strong Nash equilibrium, an α-improvement move of a coalition I strictly decreases the potential. This ensures that a sequence of improvement moves starting from the social optimum will converge because, at any time, $0 \leq \Phi \leq (1 + \ln W)cost(OPT)$.

Consider an α-improvement move of a coalition I of agents. Again, we view the move as being performed in two steps. (1) Agents $i \in I$ first drop all the edges of their strategies S_i. Let E_1 be this set of edges. (2) Agents $i \in I$ buy the edges they wish to have in their new strategies. Let E_2 be the set of edges involved. In the following, let $cost^-$ be the absolute value of the cost reduction experienced by I due to step (1). Note that $cost^-$ is equal to the cost of I in configuration $\mathcal{S}$. Let Φ^- be the absolute value of the potential drop. Similarly, let $cost^+$ be the value of the cost increase of I in step (2) and Φ^+ be the corresponding potential increase. The value of $cost^+$ is equal to the cost of I in the new configuration after strategy change. Using the definition of an α-improvement move, we find $\alpha cost^+ - cost^- < 0$. It remains to show that $cost^- \leq \Phi^-$ and $\Phi^+ \leq \alpha cost^+$, which implies $\Delta\Phi = -\Phi^- + \Phi^+ < 0$.

For any edge $e \in E$, let W_e^1 be the total weight of agents sharing e after step (1). The cost reduction experienced by I due to edge $e \in E_1$ is $cost_e^- = c(e)(W_e - W_e^1)/W_e$. For any $e \in E_1$, let Φ_e^- denote the potential drop caused by this edge. If $W_e^1 = 0$, then $cost_e^- = c(e) \leq c(e)(1 + \ln W_e) = \Phi_e^-$. If $W_e^1 > 0$, then $W_e^1 \geq 1$ and

$$cost_e^- = c(e)\frac{W_e - W_e^1}{W_e} \leq c(e)\int_{W_e^1}^{W_e} \frac{1}{z}dz = c(e)(\ln W_e - \ln W_e^1) = \Phi_e^-.$$

We conclude $cost^- = \sum_{e\in E_1} cost_e^- \leq \sum_{e\in E_1} \Phi_e^- = \Phi^-$.

For any $e \in E_2$ let W_e^2 be the total weight of agents sharing e after step (2). The cost increase experienced by I due to edge $e \in E_2$ is $cost_e^+ = c(e)(W_e^2 - W_e^1)/W_e^2$ because agents in I purchasing e have a total weight of $W_e^2 - W_e^1$. Let Φ_e^+ be the potential increase caused by $e \in E_2$. If $W_e^1 = 0$, then $\Phi_e^+ = c(e)(1 + \ln W_e^2) \leq c(e)(1 + \ln(1 + \overline{W})) \leq \alpha cost_e^+$. If $W_e^1 > 0$, then $\Phi_e^+ = c(e)(\ln W_e^2 - \ln W_e^1) = c(e)\ln(W_e^2/W_e^1)$. To establish $\Phi_e^+ \leq \alpha cost_e^+$, we prove that

$$f(W_e^2) = \ln(W_e^2/W_e^1) - (1 + \ln(1 + \overline{W}))\frac{W_e^2 - W_e^1}{W_e^2}$$

is upper bounded by 0, for all $W_e^2 \geq W_e^1$. This implies $\Phi^+ = \sum_{e\in E_2} \Phi_e^+ \leq \sum_{e\in E_2} \alpha cost_e^+ = \alpha cost^+$, because $\alpha \geq 1 + \ln(1 + \overline{W})$. Computing the first derivative of f we find that f is decreasing for values of W_e^2 between W_e^1 and $(1 + \ln(1 + \overline{W}))W_e^1$ and increasing for larger values. Since $f(W_e^1) = 0$, we obtain that f

is upper bounded by 0 for any W_e^2 with $W_e^1 \leq W_e^2 \leq (1+\ln(1+\overline{W}))W_e^1$. If $W_e^2 > (1+\ln(1+\overline{W}))W_e^1$, then $W_e^1 < W_e^2/(1+\ln(1+\overline{W}))$ and

$$(1+\ln(1+\overline{W}))(W_e^2 - W_e^1)/W_e^2 \geq \ln(1+\overline{W}).$$

Hence

$$\begin{aligned} f(W_e^2) &\leq \ln(W_e^2/W_e^1) - \ln(1+\overline{W}) \\ &= \ln(1+\frac{W_e^2 - W_e^1}{W_e^1}) - \ln(1+\overline{W}) \\ &\leq \ln(1+\overline{W}) - \ln(1+\overline{W}) \;=\; 0. \end{aligned}$$

□

Theorem 6 *In any directed or undirected graph and for any $\alpha \geq 1$, the price of anarchy of α-approximate strong Nash equilibria is upper bounded by $\frac{\alpha n}{c}(1+\ln W^c)$ if coalitions of size up to c are allowed. The price of anarchy is upper bounded by $\frac{2\alpha W}{W^c_{\max}}(1+\ln W^c_{\max})$ if coalitions of weight up to $W^c_{\max}$ are allowed.*

If there are no restrictions on the coalitions being formed and for $\alpha = 1$, we obtain the following corollary.

Corollary 2 *In any directed or undirected graph the price of anarchy of strong Nash equilibria is upper bounded by $1+\ln W$.*

Proof of Theorem 6. We generalize the proof of Theorem 2. Given an α-approximate strong Nash equilibrium $\mathcal{S} = (S_1, \ldots, S_n)$, we consider a coalition I of legal size or weight and show

$$\sum_{i \in I} cost_i(\mathcal{S}) \leq \alpha(1+\ln W_I) cost_i(OPT), \tag{6}$$

where W_I it the total weight of agents $i \in I$. If coalitions of size up to c are allowed, inequality (6) gives $\sum_{i\in I} cost_i(\mathcal{S}) \leq \alpha(1+\ln W^c)cost(OPT)$. Summing this inequality over all the $\binom{n}{c}$ coalitions of size exactly c, we obtain $cost(\mathcal{S}) \leq \frac{\alpha n}{c}(1+\ln W^c)cost(OPT)$. If coalitions of weight up to $W^c_{\max}$ are allowed, we partition the n agents into maximal possible coalitions of admissible weight. This partitioning consists of at most $2W/W^c_{\max}$ coalitions because only one of these coalitions can have a total weight of at most $W^c_{\max}/2$ and the total weight of any two coalitions is larger than $W^c_{\max}$. For each coalition of the partitioning we sum up (6) an obtain $cost(\mathcal{S}) \leq \frac{2\alpha W}{W^c_{\max}}(1+\ln W^c_{\max})cost(OPT)$.

In order to establish (6), for any fixed coalition I, we perform the same process as in the proof of Theorem 2, where sub-coalitions of I change strategy and purchase the edge set E^{OPT} of the social optimum. For any $i \in I$, let $E_i^{OPT} \subseteq E^{OPT}$ be a minimal edge set necessary to connect the terminals of agent i in the optimal solution. The process starts with $I_1 := I$. In the kth step, for $k = 1, \ldots, |I|$, agents i in the remaining sub-coalition I_k change strategies and connect their terminals using E_i^{OPT}. Since the original configuration $\mathcal{S}$ is an α-approximate strong Nash equilibrium, there must exist one agent i_k whose cost in the original configuration $\mathcal{S}$ is bounded by

$$cost_{i_k}(\mathcal{S}) \leq \alpha \sum_{e \in E_{i_k}^{OPT}} c(e) \frac{w_{i_k}}{W^e_{I_k}}, \tag{7}$$

where $W^e_{I_k}$ is the total weight of agents sharing e, i.e. $W^e_{I_k} = \sum_{i \in I_k^e} w_i$ with $I_k^e = \{i : i \in I_k \text{ and } e \in E_i^{OPT}\}$. This agent i_k leaves the process and $I_{k+1} := I_k \setminus \{i_k\}$. Summing (7) over all the $|I|$ steps, we obtain

$$\sum_{i \in I} cost_i(\mathcal{S}) \leq \alpha \sum_{k=1}^{|I|} \sum_{e \in E_{i_k}^{OPT}} c(e) \frac{w_{i_k}}{W^e_{I_k}}. \tag{8}$$

We estimate the contribution of $c(e)$, for a fixed edge $e \in E^{OPT}$, in the right-hand side expression of (8). A contribution to the sum occurs whenever an agent $i \in I$ with $e \in E_i^{OPT}$ leaves the process. Let $i_1, \ldots, i_\ell \in I$ be the agents using e, i.e. $e \in E_{i_j}^{OPT}$ for $j = 1, \ldots, \ell$, and assume that these agents are numbered according to the time when they leave the process of strategy changes. For $j = 1, \ldots, \ell$, let $s_j = w_{i_j} + \ldots + w_{i_\ell}$ be the suffix sum of these agents' weights. Then edge e contributes a total of

$$\begin{aligned}
& c(e)\left(\frac{w_{i_1}}{w_{i_1} + \ldots + w_{i_\ell}} + \frac{w_{i_2}}{w_{i_2} + \ldots + w_{i_\ell}} + \ldots + \frac{w_{i_\ell}}{w_{i_\ell}}\right) \\
\leq \;& c(e)(\sum_{j=1}^{\ell-1} \int_{s_{j+1}}^{s_j} \frac{1}{z} dz + 1) \;\leq\; c(e)(1 + \int_{s_\ell}^{s_1} \frac{1}{z} dz) \\
\leq \;& c(e)(1 + \int_{1}^{s_1} \frac{1}{z} dz) \;=\; c(e)(1 + \ln s_1) \\
\leq \;& c(e)(1 + \ln W_I).
\end{aligned}$$

The third inequality holds because $s_\ell = w_\ell \geq 1$. Summing this cost estimate over all edges $e \in E^{OPT}$, we obtain the desired bound on $\sum_{i \in I} cost_I(\mathcal{S})$. □

5.2 Lower bounds

We develop lower bounds on the performance of strong Nash equilibria in directed and undirected graphs.

Theorem 7 *In directed graphs the price of anarchy of α-approximate strong Nash equilibria is at least $\Omega(\alpha \max\{n/c, \log W\})$ if coalitions of size at most c are allowed, and at least $\Omega(\alpha \max\{W/W_{\max}^c, \log W\})$ if coalitions of weight up to $W_{\max}^c$ are allowed.*

Proof. In the setting where coalitions of size up to c are permitted, a lower bound of $\alpha n/c$ was already shown for unweighted games in Theorem 3. We first prove the lower bound of $\alpha W/W_{\max}^c$ if coalitions of weight up to $W_{\max}^c$ are feasible. Consider n agents with arbitrary weights w_i, $1 \leq i \leq n$. We use the simple network depicted in Figure 1(a) but change the costs of the two parallel edges. The expensive edge now has cost αW whereas the inexpensive one costs $W_{\max}^c + \epsilon$. Recall that all the n agents have to connect terminals s and t. The state in which all the agents establish their connection using the expensive edge forms an α-approximate strong Nash equilibrium: Any legal coalition incurs a cost of at most $\alpha W \frac{W_{\max}^c}{W} = \alpha W_{\max}^c$ on the expensive edge. Switching to the inexpensive edge results in a cost of $W_{\max}^c + \epsilon$ for the coalition, which is not attractive enough. Obviously, the social optimum routes connections via the inexpensive edge.

We next show a lower bound of $\Omega(\alpha \log W)$ for both scenarios, where either the size or the weight of a coalition is limited. W.l.o.g. let W be a power of 2 and let $n = \log_2 W + 1$. We use the graph of Figure 1(b) but change the costs of the edges. Each edge (v_i, t) now has cost α, $1 \leq i \leq n$, and edge (w, t) has cost $2 + \epsilon$. The edges (v_i, w) still have a cost of 0. Agent i, $1 \leq i < n$, has a weight of $W/2^i$ and wishes to connect terminals v_i and t. The last agent n has a weight of 1 and has to connect v_n to t. The total weight of all the n agents is exactly W. The state in which every agent i, $1 \leq i \leq n$, establishes its connection using edge (v_i, t) represents an α-approximate strong Nash equilibrium: In any coalition I of legal size or weight, the agent $i_0 \in I$ of maximum weight in I dominates the other agents in I, i.e. the weight of i_0 is at least as large as the total weight of all the other agents in I. Hence, when I changes strategy and purchases edge (w, t), agent i_0 has to pay at least $1 + \epsilon/2$, and this is not smaller than an α-fraction of the cost incurred for the private edge (v_{i_0}, t). The cost of the strong Nash equilibrium is $\alpha(1 + \log_2 W)$ while the social optimum incurs a cost of $2 + \epsilon$. □

Theorem 8 *For any $\alpha \geq 1$, there exists a family of undirected graphs, each admitting an α-approximate strong Nash equilibrium whose cost is $\Omega(\alpha\sqrt{\log W})$ times that of the social optimum.*

Proof. We extend the proof of Theorem 4 and first concentrate on $\alpha = 1$. Let W be a real weight with $\lfloor\sqrt{\log W}\rfloor \geq 2$. Again logarithms are taken to the base 3. As before we construct a graph $G = G^0$ in a recursive manner, choosing $d_{\max} = \lfloor\sqrt{\log W}\rfloor - 1$ in the case of weighted games.. In any graph G^d, $0 \leq d \leq d_{\max}$, the edge costs are the same as those defined in the proof of Theorem 4. However, the number of associated agents changes. Associated with a graph $G^{d_{\max}}$ is *one* agent of order $d_{\max}$ having a weight of $w_{d_{\max}} = W/3^{d_{\max}(d_{\max}+1)}$. Associated with a graph G^d, $0 \leq d < d_{\max}$, is *one* agent of order d having a weight of $w_d = W/3^{d(d+1)} - 3W/3^{(d+1)(d+2)}$. The total weight of all the agents is exactly W. The total weight of all the agents associated with order-d graphs G^d and the subgraphs therein is $W_d = W/3^{d^2}$, for any $0 \leq d \leq d_{\max}$.

As edge costs have not changed, the social optimum is still constant. As usual, let $\mathcal{S}$ be the configuration in which an order-d agent purchases the arc of order-d within its graph. Then $cost(\mathcal{S}) \geq d_{\max}+1 \geq \lfloor\sqrt{\log W}\rfloor$.

To show that $\mathcal{S}$ forms a strong Nash equilibrium, we can extend Lemmas 1 and 2 in a straightforward way. In the arguments agent numbers such as n_d and N_{d+1} etc. are to be replaced by weights w_d and W_{d+1}. Some of the arguments and calculations in the proofs simplify because ceilings can be ignored and fractions f and f_k^j, reflecting portions of order-d agents that defect from routing through their order-d arcs, are now equal to either 0 or 1. Finally, for $\alpha > 1$, we again scale the arc costs by α. □

6 The price of stability in undirected graphs

In this section we address the price of stability of standard Nash equilibria in weighted games. Anshelevich et al. [2] showed a lower bound of $\Omega(\log W)$ for directed graphs. Again, $W = \sum_{i=1}^{n} w_i$ is the total weight of all the agents. We prove a lower bound for undirected graphs. No super-constant lower bound was known for undirected graphs, neither for unweighted nor for weighted games.

Theorem 9 *In undirected graphs the price of stability is* $\Omega(\log W/\log\log W)$. *This lower bound holds even if each agent has to connect only a pair of terminals. Individual terminal pairs are allowed.*

Proof. We construct a family of graphs, each admitting only one Nash equilibrium. The cost of this equilibrium will be $\Omega(\log W/\log\log W)$ times that of the social optimum. The basic structure of the graphs is the same as those constructed in the proof of Theorem 4. However the parameters are chosen differently here. Let W be a positive integer with $\log W \geq 3$. Again, logarithms are taken to the base 3. Let $d_{\max} = \lfloor \log W/(\log\log W + 1)\rfloor$. Inequality $\log W \geq 3$ implies $d_{\max} \geq 1$.

In the basic graphs $G^{d_{\max}}$ a stem edge has cost $s_{d_{\max}} = 1/3^{d_{\max}}$ and the bridge of order $d_{\max}$ has cost $b_{d_{\max}} = 3/(3^{d_{\max}}\log W)$. The arc of order $d_{\max}$ costs $a_{d_{\max}} = 1/3^{d_{\max}}$. Associated with $G^{d_{\max}}$ is one order-$d_{\max}$ agent of weight $w_{d_{\max}} = W/(3\log W)^{d_{\max}}$ wishing to connect the base and the tip of $G^{d_{\max}}$.

For any d, $0 \leq d < d_{\max}$, in a graph G^d of order d, the bridge of order d has cost $b_d = 3/(3^d \log W)$ and the arc of order d costs $a_d = 1/3^d$. Associated with G^d is one order-d agent of weight $w_d = W/(3\log W)^d - 3W/(3\log W)^{d+1}$ having to connect the base and the tip of G^d. The outermost graph $G = G^0$ is the graph we will work with.

The total weight of agents associated with one order-d graph G^d and all the subgraphs therein is $W_d = W/(3\log W)^d$. This holds for $d = d_{\max}$. Suppose that the property holds for orders $d_{\max}, d_{\max}-1, \ldots, d+1$. Since a graph of order d is composed of three graphs of order $d+1$, the total weight of agents in G^d is equal to

$$W_d = w_d + 3W/(3\log W)^{d+1} = W/(3\log W)^d.$$

In particular, we obtain that the total weight of agents in $G = G^0$ is exactly W.

Proposition 3 *The least expensive path connecting the base and the tip of a graph* G^d *using only edges of* G^d *has a total edge cost of exactly* $1/3^d$, *for any* $0 \leq d \leq d_{\max}$.

Proof. The statement of the proposition holds for $d = d_{\max}$ because the arc of $G^{d_{\max}}$ has cost $a_{d_{\max}} = 1/3^{d_{\max}}$ while the path crossing the bridge has cost $b_{d_{\max}} + s_{d_{\max}} = 3/(3^{d_{\max}} \log W) + 1/3^{d_{\max}}$. Assume that the statement of the proposition holds for depths $d_{\max}, \ldots, d+1$. In G^d the arc of order d has cost $a_d = 1/3^d$ while, by induction hypothesis, any path using the bridge of order d has a cost of at least $b_d + 3 \cdot 1/3^{d+1} = 3/(3^d \log W) + 1/3^d > 1/3^d$. □

The social optimum in G buys the backbone consisting of stem edges and bridges. There exist exactly $3^{d_{\max}}$ subgraphs of order $d_{\max}$ and hence the total cost of stem edges is $3^{d_{\max}} s_{d_{\max}} = 1$. For any fixed d, $0 \leq d \leq d_{\max}$, graph $G = G^0$ contains 3^d graphs of order d, each being equipped with an order-d bridge of cost $b_d = 3/(3^d \log W)$. Thus the total cost of order-d bridges is $3^d b_d = 3^d \cdot 3/(3^d \log W) = 3/\log W$. Summing over all d we find that the total cost of bridges is upper bounded by $(d_{\max} + 1) \cdot 3/\log W \leq (2 \log W/\log\log W)(3/\log W) \leq 6/\log\log W \leq 6$. Hence the cost of the social optimum is constant.

Consider configuration $\mathcal{S}$ in which, for any graph G^d within G, the order-d agent associated with G^d purchases the order-d arc in this graph. We will prove in the following that $\mathcal{S}$ is a Nash equilibrium and that it is the only Nash equilibrium in G. As there are 3^d graphs of order d, the total cost of order-d arcs is $3^d a_d = 3^d \cdot 1/3^d = 1$ and summing over all d we obtain $cost(\mathcal{S}) = d_{\max} + 1 \geq \log W/(\log\log W + 1)$, which gives the stated lower bound on the price of stability.

In the remainder of this proof we show, in a first step, that $\mathcal{S}$ forms a Nash equilibrium and then, in a second step, that $\mathcal{S}$ is the only equilibrium in $G = G^0$.

We proceed with the proof that $\mathcal{S}$ forms an equilibrium state. Let i be an order-d agent associated with a graph G^d. We show that any strategy change performed by i yields a strictly higher cost. If i deviates from its original strategy in $\mathcal{S}$, it can establish the required connection either (1) by using a path within G^d or (2) by using a path of edges outside G^d.

In case (1), the path P_i used by agent i to connect its terminal pair has to traverse the order-d bridge in G^d, which has a cost of $b_d = 3/(3^d \log W)$. If $d = d_{\max}$, path P_i continues on the stem edge of cost $s_{d_{\max}} = 1/3^{d_{\max}}$. If $d < d_{\max}$, then P_i has to traverse three graphs of order $d+1$, the total cost of which is at least $3 \cdot 1/3^{d+1} = 1/3^d$. The total weight of agents that can share the cost of P_i is upper bounded by W_d. Thus agent i incurs a cost of at least

$$\frac{w_d}{W_d} \cdot \frac{1}{3^d}\left(1 + \frac{3}{\log W}\right).$$

If $d = d_{\max}$, then $w_d = W_d$ and the latter expression is larger than the cost of $1/3^{d_{\max}}$ incurred for buying the order-$d_{\max}$ arc in $G_{d_{\max}}$. If $d < d_{\max}$, then we have

$$\begin{aligned}
\frac{w_d}{W_d} \cdot \frac{1}{3^d}\left(1 + \frac{3}{\log W}\right) &= \frac{W/(3\log W)^d - 3W/(3\log W)^{d+1}}{W/(3\log W)^d} \cdot \frac{1}{3^d}\left(1 + \frac{3}{\log W}\right) \\
&= \left(1 - \frac{1}{\log W}\right)\left(1 + \frac{3}{\log W}\right)\frac{1}{3^d} = \frac{\log^2 W + 2\log W - 3}{\log^2 W} \cdot \frac{1}{3^d} > 1/3^d,
\end{aligned}$$

because $\log W \geq 3$. Again, buying the order-d arc of cost $a_d = 1/3^d$ is a strictly better strategy.

In case (2), we have $d > 0$ and the path P_i used by agent i crosses the bridge of order $d-1$ and visits the base of graph G^{d-1} containing G^d. The structure of P_i is depicted in Figures 3(a–c), which we used in an earlier proof; the situation is the very same here. To reach the tip of G^d, path P_i must visit the tip of G^{d-1}, from where it can continue. Path P_i must fully traverse two subgraphs of order d within G^{d-1}. Such a subgraph can be traversed on an arc of order d, having cost $a_d = 1/3^d$, where the weight of the agent that bought this arc in $\mathcal{S}$ is w_d. Thus the cost of a_d can be shared among two order-d agents. If P_i does not use the order-d arc, the total cost of edges traversing an order-d subgraph is at least $1/3^d$ and the total weight of agents sharing the edge cost is $W_d - w_d < w_d$. Thus, for the traversal of the two order-d subgraphs, agent i

pays at least

$$2\frac{w_d}{2w_d} \cdot \frac{1}{3^d} = \frac{1}{3^d}.$$

Since the traversal of the order-$(d-1)$ bridge has positive cost, path P_i incurs a cost strictly higher than that of the original strategy of i in $\mathcal{S}$.

It remains to show that $\mathcal{S}$ is the only Nash equilibrium. To this end we will prove that in any Nash equilibrium, an order-d agent associated with a given graph G^d must buy the corresponding order-d arc in G^d. In other words, an equilibrium state must be equal to $\mathcal{S}$. The desired statement that in any Nash equilibrium an order-d agent purchases the corresponding order-d arc in its graph G^d follows from the next lemma. Loosely speaking, this lemma says that in a Nash equilibrium connections in G are established locally. We first state the lemma and then explain its implications.

Lemma 3 *Consider a fixed order-d graph G^d in G and assume that in any Nash equilibrium all the agents associated with G^d and its subgraphs establish their connections using only edges of G^d. Furthermore, assume that all agents* not *associated with G^d or its subgraphs do* not *use any edges of G^d when routing their connections. Then in any Nash equilibrium the following two properties hold.*

(a) The order-d agent associated with G^d buys the arc of order d in G^d.

(b) If $d < d_{\max}$, then for any of the three order-$(d+1)$ subgraphs G_k^{d+1}, $1 \le k \le 3$, within G^d, the agents associated with G_k^{d+1} and its subgraphs establish their connections using only edges of G_k^{d+1}.

Using this lemma we can finish the proof of our theorem: For $d = 0$, trivially, all agents associated with $G = G^0$ and its subgraphs must establish connections within G^0 and there exist no agents outside G^0 that could use edges of G^0. Thus the conditions of Lemma 3 are met and we obtain that the order-0 agent buys the arc of order 0 (part (a)) and that, for any of the three subgraphs G_k^1, $1 \le k \le 3$, agents associated with any G_k^1 and its subgraphs establish connections using only edges of this graph G_k^1 (part (b)). Inductively, Lemma 3 yields that, for any d, (a) any order-d agent purchase the order-d arc within its graph and that (b) for any subgraph G^{d+1} of order $d+1$, all agents associated with G^{d+1} at its subgraphs establish the required connections locally within G^{d+1}.

Proof of Lemma 3. Part (a): Suppose that in a Nash equilibrium, an order-d agent associated with a graph G^d does not purchase the arc of order d. Let P be the path used by the agent to connect its terminal pair. Since, by assumption of the lemma, the agent establishes its connection within G^d, path P must cross the order-d bridge, see Figures 2(a) and (b). If $d = d_{\max}$, then the path traverses the stem edge of cost $s_{d_{\max}} = 1/3^{d_{\max}}$ in $G^d = G^{d_{\max}}$ to reach the tip of the graph. If $d < d_{\max}$, path P traverses the three subgraphs G_k^{d+1}, $1 \le k \le 3$, to reach the tip of G^d. When traversing the subgraphs, then by Proposition 3 path P visits edges of total cost at least $3 \cdot 1/3^{d+1} = 1/3^d$. Hence, in any case, the total cost of edges traversed by P is at least $b_d + 1/3^d = 3/(3^d \log W) + 1/3^d = (1 + 3/\log W)/3^d$. By assumption of the lemma, agents not associated with G^d or its subgraphs do not use edges of G^d. Hence the cost of P is shared by agents of total weight at most $W_d = W/(3\log W)^d$ that are associated with G^d and its subgraphs. Hence the total cost of the order-d agent is at least $\frac{w_d}{W_d}\frac{1}{3^d}(1 + \frac{3}{\log W})$. We argue that this expression is strictly larger than $a_d = 1/3^d$, which is the cost of purchasing the arc of order d. If $d = d_{\max}$, then $w_d = W_d$ and we are done. If $d < d_{\max}$, then as on the previous page we can show we have $\frac{w_d}{W_d}\frac{1}{3^d}(1 + \frac{3}{\log W}) > \frac{1}{3^d}$.

Part (b): We first prove that any order-$(d+1)$ agent associated with a graph G_k^{d+1}, $1 \le k \le 3$, establishes its connection within G_k^{d+1}. We then show that agents of order larger than $d+1$ associated with subgraphs of G_k^{d+1}, if such subgraphs exist, also route their connections within G_k^{d+1}.

In a first step we lower bound the cost incurred by order-$(d+1)$ agents if they buy edges outside their graph. In the following, if $d+1 = d_{\max}$, we set $W_{d+2} = 0$.

Claim 1 *If the agent of order $d+1$ associated with a graph G_k^{d+1} establishes its connection using edges outside G_k^{d+1}, then its total cost is at least $C = \frac{2w_{d+1}}{2w_{d+1}+9W_{d+2}} \frac{1}{3^{d+1}}(1 + \frac{3}{\log W})$.*

Proof. We first show that if the order-$(d+1)$ agent associated with G_k^{d+1} implements its connection using edges outside G_k^{d+1}, then the path P used by this agent must traverse the bridge of order d as well as the other two graphs of order $d+1$ within G^d. We consider all possible values of k and refer the reader to Figures 4(a–c) for the structure of P. The situation is the same as that described in Figures 3(a–c); the only difference is that here the outer graph has order d instead of $d-1$. Recall that strategies are simple paths connecting the desired terminals. If $k=1$ (cf. Fig. 4(a)), then the order-$(d+1)$ agent associated with G_1^{d+1} must traverse the order-d bridge of G^d. In order to reach the tip of G_1^{d+1}, path P must visit the tip of G^d, from where P traverses G_3^{d+1} and G_2^{d+1}. Similarly, if $k=2$ (cf. Fig. 4(b)), then path P traverses G_1^{d+1}, crosses the bridge of order d, travels to the tip of G^d and visits G_3^{d+1} to reach the tip of G_2^{d+1}. Finally, if $k=3$ (cf. Fig. 4(c)), path P must traverse G_2^{d+1} and G_1^{d+1}. The path then crosses the bridge of order d and travels to the tip of G^d, which is also the tip of G_3^{d+1}.

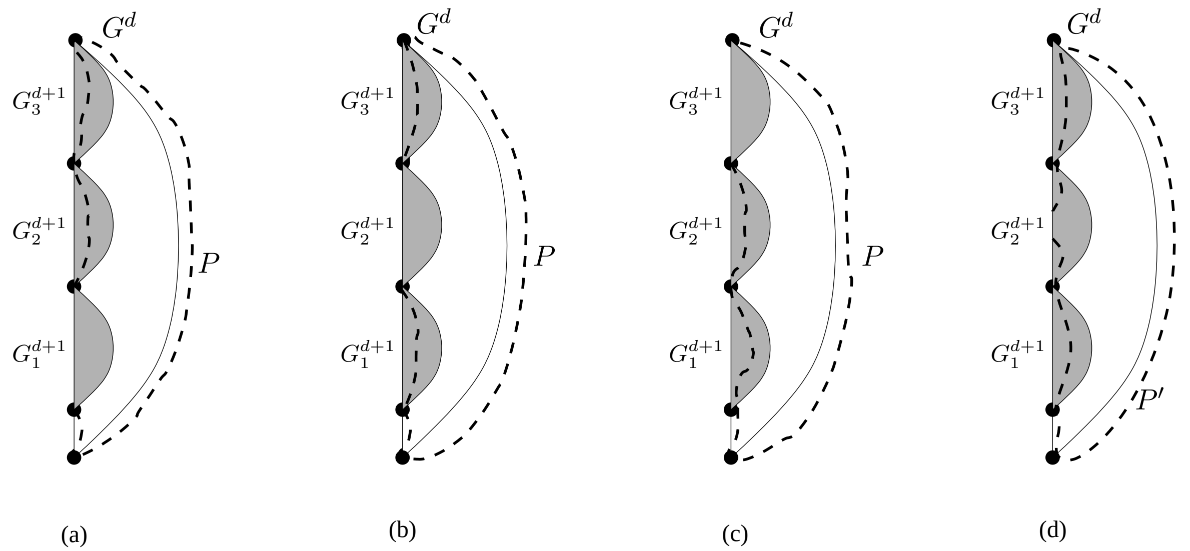

Figure 4: The paths taken by an order-$(d+1)$ agent within G_k^d, $1 \leq k \leq 3$.

We lower bound the cost incurred by the order-$(d+1)$ agent associated with G_k^{d+1} for path P. For brevity, we will denote this agent by i_k. By the assumptions of the lemma to be proven, agents not associated with G^d or its subgraphs do not use edges of G^d. Graph G^d contains three order-$(d+1)$ agents of total weight $3w_{d+1}$. If $d \leq d_{\max} - 2$, then G^d also contains 9 graphs of order $d+2$, each hosting agents of total weight W_{d+2}. We set $W_{d+2} = 0$ if $d = d_{\max} - 1$. We argued in the last paragraph that path P must cross the bridge of order d, which has a cost of $b_d = 3/(3^d \log W)$. This cost is split among agents of total weight at most $3w_{d+1} + 9W_{d+2}$. Hence, for the bridge of order d, agent i_k pay at least

$$\frac{w_{d+1}}{3w_{d+1}+9W_{d+2}} \cdot \frac{3}{3^d \log W} \geq \frac{2w_{d+1}}{2w_{d+1}+9W_{d+2}} \cdot \frac{3}{3^{d+1} \log W}. \tag{9}$$

Also, as argued in the last paragraph, path P connecting the terminals of i_k has to traverse the other two subgraphs of order $d+1$ in G^d, which are indexed $k' = k \bmod 3 + 1$ and $k'' = (k+1) \bmod 3 + 1$. To traverse one such subgraph, path P traverses edges of total cost at least $1/3^{d+1}$. We distinguish cases depending on whether the order-$(d+1)$ agents associated with $G_{k'}^{d+1}$ and $G_{k''}^{d+1}$ implement their connections using edges inside or outside their respective graphs. First, assume that the order-$(d+1)$ agents associated with $G^{d+1}_{k'}$ and $G^{d+1}_{k''}$ establish connections within their respective subgraphs. In this case path P encounters agents of total weight at most $w_{d+1} + 9W_{d+2}$ in each of these subgraphs. Thus, the traversal cost of P is shared among agents of total weight at most $2w_{d+1} + 9W_{d+2}$. We obtain that agent i_k incurs a cost of at least

$$2\frac{w_{d+1}}{2w_{d+1} + 9W_{d+2}} \cdot \frac{1}{3^{d+1}} \tag{10}$$

for the traversal of $G^{d+1}_{k'}$ and $G^{d+1}_{k''}$. Next, assume that the order-$(d+1)$ agents associated with $G^{d+1}_{k'}$ and $G^{d+1}_{k''}$ establish connections using edges outside their graphs. In this case, again, path P encounters other agents of total weight at most $w_{d+1} + 9W_{d+2}$ when traversing any of these subgraphs (in $G^{d+1}_{k'}$ the associated order-$(d+1)$ agent is not present; the analogous statement holds for $G^{d+1}_{k''}$). Hence the cost incurred in traversing any of the two subgraphs $G^{d+1}_{k'}$ and $G^{d+1}_{k''}$ is shared among agents of total weight at most $2w_{d+1} + 9W_{d+2}$ and we obtain the same cost bound as that given in (10). Finally, assume that in exactly one of the subgraphs among $G^{d+1}_{k'}$ and $G^{d+1}_{k''}$ the associated order-$(d+1)$ agent establishes its connection within their subgraph. As for the other of the two subgraphs, the associated order-$(d+1)$ agent uses edges outside its graph. W.l.o.g. let $G^{d+1}_{k'}$ be the graph where connections are made inside and let $G^{d+1}_{k''}$ be the one where connections are established using edges outside. The other case is symmetric. When P traverses $G^{d+1}_{k'}$, agents of total weight at most $2w_{d+1} + 9W_{d+2}$ are present and cost sharing on edges can be done among agents of weight at most $3w_{d+1} + 9W_{d+2}$. When P visits $G^{d+1}_{k''}$, agents of total weight at most $9W_{d+2}$ are present: The order-$(d+1)$ agent associated with $G^{d+1}_{k'}$ is not present because it uses connections inside its graph, and the order-$(d+1)$ agent associated with $G^{d+1}_{k''}$ is not present because it uses a strategy outside its graph. Hence cost sharing on edges can be done among agents of total weight at most $w_{d+1} + 9W_{d+2}$ and the cost incurred by i_k in traversing the two other order-$(d+1)$ graphs is at least

$$\left(\frac{w_{d+1}}{3w_{d+1} + 9W_{d+2}} + \frac{w_{d+1}}{w_{d+1} + 9W_{d+2}}\right)\frac{1}{3^{d+1}} \geq 2\frac{w_{d+1}}{2w_{d+1} + 9W_{d+2}} \cdot \frac{1}{3^{d+1}},$$

which is the same expression as (10). Summing the costs incurred for crossing the order-d bridge and for traversing other order-$(d+1)$ graphs, see (9) and (10), we conclude that agent i_k pays at least the cost of C stated in the claim. □

If the order-$(d+1)$ agent associated with G^{d+1}_k purchases the arc of order $d+1$ within G^{d+1}_k, its cost is at most $a_{d+1} = 1/3^{d+1}$. A strategy using edges outside the graph incurs a costs of at least C as stated in the above claim. We show that $C > a_{d+1}$, which proves that in a Nash equilibrium the order-$(d+1)$ agent associated with G^{d+1}_k establishes the required connection via the order-$(d+1)$ arc. If $d = d_{\max} - 1$, then $W_{d+2} = 0$ and we are done because $C = \frac{1}{3^{d+1}}(1 + \frac{3}{\log W}) > \frac{1}{3^{d+1}}$.

If $d \leq d_{\max} - 2$, then

$$C = \frac{2w_{d+1}}{2w_{d+1} + 9W_{d+2}}\frac{1}{3^{d+1}}(1 + \frac{3}{\log W}) = \left(1/\left(1 + \frac{9W_{d+2}}{2w_{d+1}}\right)\right)\frac{1}{3^{d+1}}(1 + \frac{3}{\log W}).$$

It remains to show $9W_{d+2}/(2w_{d+1}) < 3/\log W$, which proves the desired inequality $C > 1/3^{d+1}$. We have

$$\frac{9W_{d+2}}{2w_{d+1}} = \frac{9W/(3\log W)^{d+2}}{2W/(3\log W)^{d+1} - 6W/(3\log W)^{d+2}} = \frac{3}{2\log W - 2} \leq \frac{3}{\log W}.$$

The last inequality holds because $\log W \geq 3$.

To finish the proof of part (b) of the lemma we have to show that if $d \leq d_{\max} - 2$, then any agent i of order $d+2$ or larger that is associated with a subgraph of G^{d+1}_k, $1 \leq k \leq 3$, establishes its connection within G^{d+1}_k. Suppose this were not the case. Then agent i chooses a path P that leaves G^{d+1}_k through its base. Figure 4(d) shows a sample path for $k = 2$. To connect to the desired terminal, path P must visit the tip of G_k^{d+1} from where it can continue on edges inside G_k^{d+1}. Since P uses edges outside G_k^{d+1}, it does not use the arc of order $d+1$ in G_k^{d+1} and hence must traverse the bridge of order d in G^d. The cost of this bridge is shared by agents of total weight at most $9W_{d+2}$ because we have shown that all the agents of order $d+1$ in G^d establish connections within their respective subgraphs and the order-d agent associated with G^d purchases the arc of order d (see part(a)). Let P' be the subpath of P connecting the base and the tip of G_k^{d+1}. On P' agent i incurs a cost of at least

$$cost(P') \geq \frac{w(i)}{9W_{d+2}} \frac{3}{3^d \log W},$$

where $w(i)$ is the weight of agent i. In the given Nash equilibrium, consider the strategy used by the order-$(d+1)$ agent associated with G_k^{d+1}. The cost of this agent is at most $1/3^{d+1}$ because this is the cost incurred when buying the order-$(d+1)$ arc in G_k^{d+1}, which is always an option. This order-$(d+1)$ agent has to connect the base and the tip of G_k^{d+1} and, as shown in the previous paragraphs, use edges within G_k^{d+1}. Now, agent i can replace P' by the strategy used by the order-$(d+1)$ agent of G_k^{d+1}, incurring a cost of at most $w(i)/(w(i)+w_{d+1}) \cdot 1/3^{d+1} \leq w(i)/w_{d+1} \cdot 1/3^{d+1}$. We show that the latter expression is strictly smaller than the cost $cost(P')$, contradicting the fact that the configuration in which i used edges outside G_k^{d+1} was a Nash equilibrium. Inequality $cost(P') > w(i)/w_{d+1} \cdot 1/3^{d+1}$ is equivalent to showing $9W_{d+2}/w_{d+1} < 9/\log W$. We have

$$\frac{9W_{d+2}}{w_{d+1}} = \frac{9W/(3\log W)^{d+2}}{W/(3\log W)^{d+1} - 3W/(3\log W)^{d+2}} \leq \frac{3}{\log W - 1} < \frac{9}{\log W},$$

since $\log W \geq 3$. □

This completes the proof of the theorem. □

The lower bound of Theorem 9 is nearly tight. Firstly, the potential function arguments of the proof of Theorem 5 imply that there exists an α-approximate Nash equilibrium whose cost is at most $1 + \ln W$ times that of the social optimum if $\alpha \geq 1 + \ln(1 + w_{\max})$. Here $w_{\max}$ is the maximum weight of any agent. Secondly, Chen and Roughgarden [7] showed that in directed graphs, for any $\alpha = \Omega(\log w_{\max})$, the price of stability of α-approximate Nash equlibria is $O((\log W)/\alpha)$. This result can be extended to undirected graphs.

7 Conclusions

In this paper we have investigated the value of coordination in network design games. We have developed lower and upper bounds on the price of anarchy attained by strong Nash equilibria in unweighted and weighted games, considering both undirected and directed graphs. It shows that strong Nash equilibria achieve much better performance ratios than standard Nash equilibria and that these ratios are often as good as those of the best standard equilibrium states. There is still room for improvements. For undirected graphs we have developed an upper bound of $H_n \approx \ln n$ and a lower bound of $\Omega(\sqrt{\log n})$ on the price of anarchy in unweighted games. In weighted games the bounds are $1 + \ln W$ and $\Omega(\sqrt{\log W})$, respectively. An interesting open problem is to determine the true ratios for undirected graphs.

Furthermore, in this paper we have also devised the first super-constant lower bound on the price of stability in unweighted graphs. More specifically we proved a lower bound of $\Omega(\log W/\log\log W)$ for weighted network design games. A challenging open problem is to determine the price of stability in unweighted games.

8 Acknowledgements

We thank Avrim Blum, Anupam Gupta and Katrina Ligett for many interesting technical discussions.

References

[1] N. Andelman, M. Feldman and Y. Mansour. Strong price of anarchy. *Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2007.

[2] 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. *Proc. 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 295–304, 2004.

[3] E. Anshelevich, A. Dasgupta, E. Tardos and T. Wexler. Near optimal network design with selfish agents. *Proc. 35th Annual ACM Symposium on Theory of Computing (STOC)*, 511–520, 2003.

[4] R.J Aumann. Acceptable points in general cooperative n-person games. In A.W. Tucker and R.D. Luce (eds.), *Contributions to the Theory of Games, Vol. IV, Annals of Mathematics Studies*, 40:287–324, 1959.

[5] V. Bala and S. Goyal. A non-cooperative model of network formation. *Econometrica*, 68:1181–1229, 2000.

[6] C. Chekuri, J. Chuzhoy, L. Lewin-Eytan, J. Naor and A. Orda. Non-cooperative multicast and facility location games. *Proc. 7th ACM Conference on Electronic Commerce (EC)*, 72–81, 2006.

[7] H.-L. Chen and T. Roughgarden. Network design with weighted players. *Proc. 18th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, 29–38, 2006.

[8] J. Corbo and D. Parkes. The price of selfish behavior in bilateral network formation. *Proc. 24th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 99–107, 2005.

[9] N.R. Devanur, M. Mihail and V.V. Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Systems*, 39:11-22, 2005.

[10] A. Epstein, M. Feldman and Y. Mansour. Strong equilibrium in cost sharing connection games. *Proc. 8th ACM Conference on Electronic Commerce*, 84–92, 2007.

[11] A. Fabrikant, A. Luthra, E. Maneva, C.H. Papadimitriou and S. Shenker. On a network creation game. *Proc. 22nd Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 347–351, 2003.

[12] J. Feigenbaum, C.H. Papadimitriou and S. Shenker. Sharing the cost of multicast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001.

[13] A. Fiat, H. Kaplan, M. Levy, S. Olonetsky and R. Shabo. On the price of stability for designing undirected networks with fair cost allocations. *Proc. 33rd International Colloquium on Automata, Languages and Programming (ICALP)*, Springer LNCS 4051, 608–618, 2006.

[14] H. Haller and S. Sarangi. Nash networks with heterogeneous agents. Working paper, Department of Economics, Louisiana State University, no. 2003-06, 2003.

[15] S. Herzog, S. Shenker and D. Estrin. Sharing the "cost" of multicast trees: an axiomatic analysis. *IEEE/ACM Transactions on Networking*, 5:847–860, 1997.

[16] M. Hoefer. Non-cooperative tree creation. *Proc. 31st International Symposium on Mathematical Foundations of Computer Science (MFCS)*, Springer LNCS 4162, 517–527, 2006.

[17] M. Jackson and A. van den Nouweland. Strongly stable networks. *Games and Economic Behavior*, 51:420–444, 2005.

[18] K. Jain and V. Vazirani. Applications of approximation algorithms to cooperative games. *Proc. 33rd Annual ACM Symposium on Theory of Computing (STOC)*, 364–372, 2001.

[19] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Proc. 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, Springer LNCS 1564, 404–413, 1999.

[20] D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[21] A. van den Nouweland. Models of network formation in cooperative games. In *Group Formation in Economics; Networks, Clubs, and Coalitions*, G. Demange and M. Wooders (eds), Cambridge University Press, 58–88, 2005.

[22] M. Pal and E. Tardos. Group strategyproof mechanisms via primal-dual algorithms. *Proc. 44th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 584-593, 2003.

[23] A. Vetta. Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions. *Proc. 43rd Symposium on the Foundations of Computer Science (FOCS)*, 416–425, 2002.

Appendix

Proof of Proposition 1. We prove the result for undirected graphs and then show how to direct edges to obtain the desired statement for directed networks as well. Consider the graph given in Figure 5. We have a vertex set $V = \{v_1, v_2, v_3, w_1, w_2, w_3, t\}$, where vertex w_i is connected to t via a *main edge* $\{w_i, t\}$ of cost 1, $1 \leq i \leq 3$. Furthermore, there are *auxiliary edges* $\{v_i, w_i\}$ of cost 1/2 and *auxiliary edges* $\{v_i, w_{i \bmod 3+1}\}$ of cost $1/2 + \epsilon$, $1 \leq i \leq 3$. Here $\epsilon > 0$ is an arbitrarily small value. Associated with the graph are three agents, where agent i has to connect terminals v_i and t, $1 \leq i \leq 3$. We will consider all possible states and show that none represents a strong Nash equilibrium. Any state in which all of the three main edges are purchased does

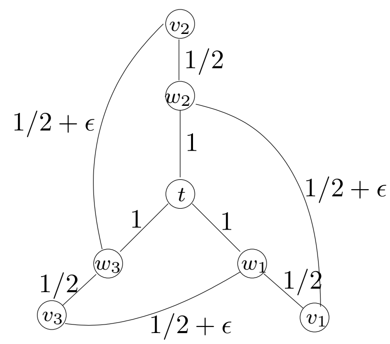

Figure 5: A graph without a strong Nash equilibrium.

not form a strong Nash equilibrium because two agents i and $i' = i \bmod 3 + 1$ could team up, sharing main edge $\{w_{i'}, t\}$. As the original cost of each of the two agents was at least $1 + 1/6 = 7/6$ and the new cost is at most $1 + \epsilon$, this yields a cost reduction for each member of the coalition.

Next suppose that there exists a strong Nash equilibrium in which two agents share a main edge $\{w_i, t\}$ while the third agent buys a second main edge $\{w_j, t\}$, $j \neq i$. Then agent i is one of the agents sharing $\{w_i, t\}$ and connects to w_i using edge $\{v_i, w_i\}$ since otherwise agent i could strictly improve its cost by purchasing a third of $\{w_i, t\}$. We now distinguish cases depending on whether $j = i \bmod 3 + 1$ or $j = (i+1) \bmod 3 + 1$.

If $j = i \bmod 3 + 1$, then agent j must be the one buying $\{w_j, t\}$ as connecting v_j to w_i requires the traversal of three auxiliary edges the cost of which is at least $1/2 + \epsilon/3$. This cost is higher than that of buying $\{v_j, w_j\}$ and hence agent j would prefer to share $\{w_j, t\}$ instead of $\{w_i, t\}$. Thus agent $i'' = (i+1) \bmod 3+1$ shares main edge $\{w_i, t\}$ and connects to w_i using edge $\{v_{i''}, w_i\}$ of cost $1/2 + \epsilon$ because any other path of auxiliary edges has a strictly higher cost. We conclude that agent j pays a cost of $3/2$ and agent i'' a cost of $1/2 + \epsilon + 1/2 = 1 + \epsilon$. Now agents j and i'' can form a coalition, sharing main edge $\{w_{i''}, t\}$. The new cost of agent j is $1/2 + \epsilon + 1/2 < 3/2$ and the new cost of agent i'' is $1/2 + 1/2 < 1 + \epsilon$, contradicting the assumption that the original configuration was a strong Nash equilibrium.

If $j = (i+1) \bmod 3 + 1$, then again agent j buys main edge $\{w_j, t\}$: If agent j shared $\{w_i, t\}$ and agent $i' = i \bmod 3+1$ bought $\{w_j, t\}$, agent i' would connect to w_j using edge $\{v_{i'}, w_j\}$ and agent j would connect to w_i using edge $\{v_j, w_i\}$ as other paths of auxiliary edges are strictly more expensive. Both agents pay a cost of $1/2 + \epsilon$ for these connections. In this situation agent j could strictly improve its cost by connecting to w_j and sharing $\{w_j, t\}$ instead of $\{w_i, t\}$. We conclude that agent i' shares main edge $\{w_i, t\}$ and connects to w_i at a cost of $1 + \epsilon + 1/4$. Now agent i' can improve its cost by buying edge $\{v_{i'}, w_j\}$ at a cost of $1/2 + \epsilon$ and sharing edge $\{w_j, t\}$ instead of $\{w_i, t\}$. We obtain a contradiction to the fact that the original configuration was a strong Nash equilibrium.

We finally have to investigate the case that a configuration buys only one main edge $\{w_i, t\}$, the cost of which is shared among the three agents. Then agent $i'' = (i+1) \bmod 3 + 1$ connects to w_i using edge $\{v_{i''}, w_i\}$ and agent $i' = i \bmod 3 + 1$ connects to w_i using a path of auxiliary edges that results in a cost of at least $1 + \epsilon + 1/4$. Hence the total cost of i' is at least $1 + 1/4 + 1/3 + \epsilon > 3/2$ and agent i' can improve its strategy by buying edges $\{v_{i'}, w_{i'}\}$ and $\{w_{i'}, t\}$.

We note that the graph can be extended to any agent number n by inserting nodes $v_4, \ldots, v_n$ affiliated with agents numbered 4 to n, where agent i wishes to connect v_i to t, $4 \leq i \leq n$. Each such v_i is connected to t via a private edge.

This concludes the analysis of undirected graphs. To obtain the result for directed graphs we simply direct edges towards the destination t. We have *main edge* (w_i, t) as well as *auxiliary edges* (v_i, w_i) and $(v_i, w_{i \bmod 3+1})$, $1 \leq i \leq 3$. Directing the edges only restricts the set of possible states, while all strategy changes proposed above can still be performed. □